

Causal Theories: A Categorical Perspective on Bayesian Networks



Brendan Fong
Wolfson College
University of Oxford

A thesis submitted for the degree of
*Master of Science in Mathematics and the Foundations of Computer
Science*

Trinity 2012

Acknowledgements

It's been an amazing year, and I've had a good time learning and thinking about the contents of this essay. A number of people have had significant causal influence on this.

Foremost among these is my dissertation supervisor Jamie Vicary, who has been an excellent guide throughout, patient as I've jumped from idea to idea and with my vague questions, and yet careful to ensure I've stayed on track. We've had some great discussions too, and I thank him for them. John Baez got me started on this general topic, has responded enthusiastically and generously to probably too many questions, and, with the support of the Centre for Quantum Technologies, Singapore, let me come visit him to pester him with more. Bob Coecke has been a wonderful and generous general supervisor, always willing to talk and advise, and has provided many of the ideas that lurk in the background of those here. I thank both of them too. I also thank Rob Spekkens, Dusko Pavlovic, Prakash Panangaden, and Samson Abramsky for some interesting discussions and detailed responses to my queries.

For help of perhaps a less technical nature, but still very much appreciated, I thank Kati, Ross, and the rest of the MFoCS crowd, as well as Fiona, Jono, Daniel, and Shiori, for being happy to listen to me rant confusedly, eat cake, and generally making life better. And finally thank you to my brothers and my parents for their constant support and understanding, even as I've been far away.

Abstract

In this dissertation we develop a new formal graphical framework for causal reasoning. Starting with a review of monoidal categories and their associated graphical languages, we then revisit probability theory from a categorical perspective and introduce Bayesian networks, an existing structure for describing causal relationships. Motivated by these, we propose a new algebraic structure, which we term a causal theory. These take the form of a symmetric monoidal category, with the objects representing variables and morphisms ways of deducing information about one variable from another. A major advantage of reasoning with these structures is that the resulting graphical representations of morphisms match well with intuitions for flows of information between these variables. These categories can then be modelled in other categories, providing concrete interpretations for the variables and morphisms. In particular, we shall see that models in the category of measurable spaces and stochastic maps provide a slight generalisation of Bayesian networks, and naturally form a category themselves. We conclude with a discussion of this category, classifying the morphisms and discussing some basic universal constructions.

Contents

Introduction	1
1 Preliminaries on Monoidal Categories	4
1.1 Monoidal categories	5
1.2 Graphical calculi	10
1.3 Example: the theory of monoids	14
2 Categorical Probability Theory	17
2.1 The category of measurable spaces	17
2.2 Measures and integration	20
2.3 The category of stochastic maps	23
2.4 Deterministic stochastic maps	27
2.5 Aside: the Giry monad	30
3 Bayesian Networks	32
3.1 Conditionals and independence	32
3.2 Bayesian networks	35
4 Causal Theories	41
4.1 The category associated to a causal structure	41
4.2 Interpretations of causal theories	46
4.3 Application: visualising Simpson’s paradox	54
5 The Structure of Stochastic Causal Models	59
5.1 Morphisms of stochastic causal models	59
5.2 Basic constructions in categories of stochastic causal models	63
Further Directions	68
Bibliography	70

Introduction

From riding a bicycle to buying flowers for a friend, causal relationships form a basic and ubiquitous framework informing, at least informally, how we organise, reason about, and choose to interact with the world. It is perhaps surprising then that ideas of causality often are entirely absent from our formal scientific frameworks, whether they directly be models of the world, such as theories of physics, or methods for extracting information from data, as in the case of statistical techniques. It is the belief of the author that there remains much to be gained from formalising our intuitions regarding causality.

Indeed, taking the view that causal relationships are fundamental physical facts about the world, it is interesting to discover and discuss these in their own right. Even if one views causality only as a convenient way to organise information about dependencies between variables, however, it is hard to see how introducing such notions into formal theories, rather than simply ignoring these intuitions, will not benefit at least some of them. The artificial intelligence community gives a tangible example of this, with the widespread use of Bayesian networks indicating that causal relationships provide a far more efficient way to encode, update, and reason with information about random variables than simply working with the entire joint variable.

In what follows we lay out the beginnings of a formal framework for reasoning about causality. We shall do this by extending the aforementioned existing ideas for describing causal relationships between random variables through the use of category theory, and in particular the theory of monoidal categories.

Overview of the literature

More precisely, in this dissertation we aim to bridge three distinct ideas. The first is the understanding of probability theory and probabilistic processes from a categorical perspective. For this we work with a category first defined by Lawvere a half-century ago in the unpublished manuscript [13], in which the objects are sets equipped with a σ -algebra and the morphisms specify a measure of the codomain for each element

of the domain, subject to a regularity condition. These ideas were later developed, in 1982, in a short paper by Giry [9], and have been further explored by Doberkat [6], Panangaden [15], and Wendt [25], among others, in recent years.

The second idea is that of using graphical models to depict causal relationships. Termed Bayesian networks, these were first discussed predominantly in the machine learning community in the 1980s, with the seminal work coming in the influential book Pearl [18]. Since then Bayesian networks have been used extensively to discuss causality from both computational and philosophical perspectives, as can be seen in recent books Pearl [19] and Williamson [26].

The third body of work, which will serve as a framework to unite the above two ideas, is the theory of monoidal categories and their graphical calculi. An introductory exposition of monoidal categories can be found in Mac Lane [14], while the survey by Selinger [20] provides an excellent overview of the graphical ideas. Here our work is in particular influenced by that of the very recent paper by Coecke and Spekkens [4], which uses monoidal categories to picture Bayesian inference, realising Bayesian inversion as a compact structure on an appropriate category.

Outline

Since they serve as the underlying framework for this thesis, we begin with a chapter reviewing the theory of monoidal categories, the last of the above ideas. We conclude this first chapter by discussing how the idea of a monoid can be generalised through a category we call the ‘theory of monoids’, with monoids themselves being realised as monoidal functors from this category into the category **Set** of sets and functions, while ‘generalised monoids’ take the form of monoidal functors from the theory of monoids into other categories. It is in this sense the ‘causal theories’ that we will define are theories. In Chapter 2 we turn our attention to reviewing the basic ideas of measure theoretic probability theory from a categorical viewpoint. Here we pay particular attention to the category, which we shall call **Stoch**, defined by Lawvere. Chapter 3 then provides some background on Bayesian networks, stating a few results about how they capture causal relationships between random variables through ideas of conditional independence. This motivates the definition of a causal theory, which we present in Chapter 4. Following our exploration of these categories and how to represent their morphisms graphically, we turn our attention to their models. Although models in **Set** and **Rel** are interesting, we spend most of the time discussing models in Lawvere’s category **Stoch**. This chapter is concluded by a discussion of confounding variables and Simpson’s paradox, where we see some of the strengths of

causal theories and in particular their graphical languages. In short, we show we can take the directed graph structure of a Bayesian network more seriously than just a suggestive depiction of dependencies. In the final chapter, Chapter 5, we discuss some properties of the category of stochastic causal models of a fixed causal theory. These are models in a certain full subcategory of **Stoch** that omits various pathological eventuations, and have a close relationship with Bayesian networks.

New contributions

The main contribution of this dissertation is the presentation of a new algebraic structure: causal theories. These are a type of symmetric monoidal category, and we will discuss how these capture the notion of causal relationships between variables; deterministic, possibilistic and probabilistic models of these categories; how their graphical calculi provide intuitive representations of reasoning and information flow; and some of the structure of the category of probabilistic models. In doing so we also move the discussion of Bayesian networks from the finite setting preferred by the computationally-focussed Bayesian network community to a more general setting capable of handling non-discrete probability spaces.

In particular, I claim all results of Chapters 4 and 5 as my own as well as, except for Proposition 2.17, the discussion of deterministic stochastic maps of Section 2.4.

Chapter 1

Preliminaries on Monoidal Categories

Our aim is to explore representations of causal relationships between random variables from a categorical perspective. In this chapter we lay the foundations for this by introducing the basic language we will be working with—the language of monoidal categories—and, by way of example, informally discussing the notion of a ‘theory’.

We begin with a review of the relevant notions from category theory. Recall that a *category* \mathcal{C} consists of a collection $\text{Ob}\mathcal{C}$ of *objects*, for each pair A, B of objects a set $\text{Mor}(A, B)$ of *morphisms*, and for each triple A, B, C of objects a function, or *composition rule*, $\text{Mor}(A, B) \times \text{Mor}(B, C) \rightarrow \text{Mor}(A, C)$, such that the composition rule is associative and obeys a unit law. We shall write $A \in \mathcal{C}$ if A is an object of the category \mathcal{C} , and f in \mathcal{C} if f is a morphism in the category \mathcal{C} . As we shall think of them, categories are the basic algebra structure capturing the idea of composable processes, with the objects of a category different systems of a given type, and the morphisms processes transforming one system into another.

We further remind ourselves that a *functor* is a map from one category to another preserving the composition rule, that under a mild size constraint the collection of categories itself forms a category with functors as morphisms, and that in this category products exist. Moreover, the set of functors between any two categories itself has the structure of a category in a standard, nontrivial way, and we call the morphisms in this category *natural transformations*, with the invertible ones further called *natural isomorphisms*. Two categories are *equivalent* if there exists a functor in each direction between the two such that their compositions in both orders are naturally isomorphic to the identity functor.

The reader seeking more detail is referred to Mac Lane [14], in particular Chapters I and II. In general our terminology and notation for categories will follow the conventions set out there.

1.1 Monoidal categories

The key structure of interest to us in the following is that of a symmetric monoidal category. A monoidal category is a category with two notions of composition—ordinary categorical composition and the monoidal composition—, and symmetric monoidal categories may be thought of as the algebraic structure of processes that may occur simultaneously as well as sequentially. These categories are of special interest as they may be described precisely with a graphical notation possessing a logic that agrees well with natural topological intuitions. Among other things, this has been used to great effect in describing quantum protocols by Abramsky and Coecke [2], and this work in part motivates that presented here.

Definition 1.1 (Monoidal category). A *monoidal category* $(\mathcal{C}, \otimes, I, \alpha, \rho, \lambda)$ consists of a category \mathcal{C} , together with a functor $\otimes : \mathcal{C} \times \mathcal{C} \rightarrow \mathcal{C}$, a distinguished object $I \in \mathcal{C}$, for all objects $A, B, C \in \mathcal{C}$ isomorphisms $\alpha_{A,B,C} : (A \otimes B) \otimes C \rightarrow A \otimes (B \otimes C)$ in \mathcal{C} natural in A, B, C , and for all objects $A \in \mathcal{C}$ isomorphisms $\rho_A : A \otimes I \rightarrow A$ and $\lambda_A : I \otimes A \rightarrow A$ in \mathcal{C} natural in A . To form a monoidal category, this data is subject to two equations: the *pentagon equation*

$$\begin{array}{ccccc}
 & & ((A \otimes B) \otimes C) \otimes D & & \\
 & \swarrow^{\alpha_{A,B,C} \otimes \text{id}_D} & & \searrow^{\alpha_{(A \otimes B), C, D}} & \\
 (A \otimes (B \otimes C)) \otimes D & & & & (A \otimes B) \otimes (C \otimes D) \\
 \downarrow^{\alpha_{A, (B \otimes C), D}} & & & & \swarrow^{\alpha_{A, B, (C \otimes D)}} \\
 A \otimes ((B \otimes C) \otimes D) & & & & \\
 \searrow^{\text{id}_A \otimes \alpha_{B, C, D}} & & & \swarrow^{\alpha_{A, B, (C \otimes D)}} & \\
 & & A \otimes (B \otimes (C \otimes D)) & &
 \end{array}$$

and the *triangle equation*

$$\begin{array}{ccc}
 (A \otimes I) \otimes B & \xrightarrow{\alpha_{A,I,B}} & A \otimes (I \otimes B) \\
 \rho_A \otimes \text{id}_B \searrow & & \swarrow \text{id}_A \otimes \lambda_B \\
 & A \otimes B &
 \end{array}$$

We call \otimes the *monoidal product*, I the *monoidal unit*, the isomorphisms α, α^{-1} *associators*, and the isomorphisms ρ, ρ^{-1} and λ, λ^{-1} *right-* and *left-unitors* respectively. Collectively, we call the associators and unitors the *structure maps* of our monoidal category. We will often just write \mathcal{C} for a monoidal category $(\mathcal{C}, \otimes, I, \alpha, \rho, \lambda)$, leaving the remaining data implicit.

The associators express the fact that the product objects $(A \otimes B) \otimes C$ and $A \otimes (B \otimes C)$ are in some sense the same—they are isomorphic via some canonical isomorphism—, while the unitors express the fact that $A, A \otimes I$, and $I \otimes A$ are the same. If these objects are in fact equal, and the structure maps are simply identity maps, then we say that our monoidal category is a *strict monoidal category*. In this case then any two objects that can be related by structure maps are equal, and so we may write objects without parentheses and units without ambiguity. Although, importantly, this is not true in all cases, it is essentially true: loosely speaking, the triangle and pentagon equations in fact imply that any diagram of their general kind, expressing composites of structure maps between different ways of forming the monoidal product of some objects, commutes. This is known as Mac Lane’s coherence theorem for monoidal categorise; see Mac Lane [14, Corollary of Theorem VII.2.1] for a precise statement and proof.

In a monoidal category, the objects $A \otimes B$ and $B \otimes A$ need not in general be related in any way. In the cases we will interest ourselves, however, we will not want the order in which we write the objects in a tensor product to matter—all products consisting of a given collection of objects should be isomorphic, and isomorphic in a way we need not worry about the isomorphism itself. This additional requirement turns a monoidal category into a symmetric monoidal category.

Definition 1.2 (Symmetric monoidal category). A *symmetric monoidal category* $(\mathcal{C}, \otimes, I, \alpha, \rho, \lambda, \sigma)$ consists of a monoidal category $(\mathcal{C}, \otimes, I, \alpha, \rho, \lambda)$ together with a collection of isomorphisms $\sigma_{A,B} : A \otimes B \rightarrow B \otimes A$ natural in A and B such that

$\sigma_{B,A} \circ \sigma_{A,B} = \text{id}_{A \otimes B}$ and such that for all objects A, B, C the hexagon

$$\begin{array}{ccc}
 & (A \otimes B) \otimes C & \\
 \alpha_{A,B,C} \swarrow & & \searrow \sigma_{A,B} \otimes \text{id}_C \\
 A \otimes (B \otimes C) & & (B \otimes A) \otimes C \\
 \sigma_{A,B} \otimes \text{id}_C \downarrow & & \downarrow \alpha_{B,A,C} \\
 (B \otimes C) \otimes A & & B \otimes (A \otimes C) \\
 \alpha_{B,C,A} \searrow & & \swarrow \text{id}_B \otimes \sigma_{A,C} \\
 & B \otimes (C \otimes A) &
 \end{array}$$

commutes.

We call the isomorphisms $\sigma_{-, -}$ *swaps*.

As for monoidal categories, we have a coherence theorem for symmetric monoidal categories, stating in essence that all diagrams composed of identities, associators, unitors, and swaps commute. Details can again be found in Mac Lane [14, Theorem XI.1.1].

Examples 1.3 (FVect, Mat). An historically important example of a symmetric monoidal category is that of $\mathbf{FVect}_{\mathbb{R}}$, the category of finite vector spaces over \mathbb{R} with linear maps as morphisms, tensor product as monoidal product. Here \mathbb{R} is the monoidal unit, and the structure maps are the obvious isomorphisms of tensor products of vector spaces. Note that this is not a strict symmetric monoidal category: it is not true for real vector spaces U, V, W that we consider $(U \otimes V) \otimes W$ and $U \otimes (V \otimes W)$ as equal, but we do always have a canonical isomorphism between the two.

A related strict symmetric monoidal category is $\mathbf{Mat}(\mathbb{R})$, the category with objects natural numbers, morphisms from $m \in \mathbb{N}$ to $n \in \mathbb{N}$ given by $n \times m$ matrices over \mathbb{R} , composition given by composition of matrices, monoidal product given by multiplication on objects and Kronecker product of matrices on morphisms.

Example 1.4 (Set). The category \mathbf{Set} of sets and functions forms a symmetric monoidal category with the cartesian product \times . In this category any singleton set $\{*\}$ may be taken as the monoidal unit. Indeed, any category with finite products can be viewed as a symmetric monoidal category by taking the binary categorical product as the monoidal product, and the terminal object as the monoidal unit. The associators, unitors, and swaps are then specified by the unique isomorphisms given by the universal property of the product.

Example 1.5 (Rel). The category **Rel** of sets and relations forms a symmetric monoidal category with cartesian product \times and unit $\{*\}$. Here the monoidal product $r \otimes s : X \times Z \rightarrow Y \times W$ of relations $r : X \rightarrow Y$ and $s : Z \rightarrow W$ is the relation such that $(x, z) \in X \times Z$ is related to $(y, w) \in Y \times W$ if and only if x is related to y by r and z is related to w by s .

Intuitively, the standard embedding of **Set** into **Rel**, given by viewing functions as relations, is an embedding that respects the monoidal structure. To make this precise we need to talk about monoidal functors.

When working with monoidal categories, it is often desirable to have functors between these categories preserve the monoidal structure, and to have natural transformations between these functors preserve the monoidal structure too. The same is true in the case of functors between symmetric monoidal categories. We thus introduce the notions of monoidal functors, symmetric monoidal functors, and monoidal natural transformations.

Definition 1.6 (Monoidal functor). Let $\mathcal{C}, \mathcal{C}'$ be monoidal categories. A *monoidal functor* $(F, F_\otimes, F_*) : \mathcal{C} \rightarrow \mathcal{C}'$ from \mathcal{C} to \mathcal{C}' consists of a functor $F : \mathcal{C} \rightarrow \mathcal{C}'$, for all objects $A, B \in \mathcal{C}$ morphisms

$$F_{\otimes, A, B} : F(A) \otimes F(B) \rightarrow F(A \otimes B)$$

in \mathcal{C}' which are natural in A and B , and for the units I of \mathcal{C} and I' of \mathcal{C}' a morphism $F_* : I' \rightarrow F(I)$ in \mathcal{C}' , such that for all $A, B, C \in \mathcal{C}$ the hexagon

$$\begin{array}{ccc}
 & (FA \otimes FB) \otimes FC & \\
 \alpha_{FA, FB, FC} \swarrow & & \searrow F_{\otimes, A, B} \otimes \text{id}_{FC} \\
 FA \otimes (FB \otimes FC) & & F(A \otimes B) \otimes FC \\
 \text{id}_{FA} \otimes F_{\otimes, B, C} \downarrow & & \downarrow F_{\otimes, A \otimes B, C} \\
 FA \otimes F(B \otimes C) & & F((A \otimes B) \otimes C) \\
 F_{\otimes, A, B \otimes C} \searrow & & \swarrow F\alpha_{A, B, C} \\
 & F(A \otimes (B \otimes C)) &
 \end{array}$$

and the two squares

$$\begin{array}{ccc}
F(A) \otimes I' & \xrightarrow{\rho} & F(A) \\
\text{id} \otimes F_* \downarrow & & \uparrow F\rho \\
F(A) \otimes F(I) & \xrightarrow{F_{\otimes, A, I}} & F(A \otimes I)
\end{array}
\qquad
\begin{array}{ccc}
I' \otimes F(A) & \xrightarrow{\lambda} & F(A) \\
F_* \otimes \text{id} \downarrow & & \uparrow F\lambda \\
F(I) \otimes F(A) & \xrightarrow{F_{\otimes, I, A}} & F(I \otimes A)
\end{array}$$

commute.

We further say a monoidal functor is a *strong monoidal functor* if the morphisms $F_{\otimes, A, B}$ and F_* are isomorphisms for all $A, B \in \mathcal{C}$.

Definition 1.7 (Symmetric monoidal functor). A symmetric monoidal functor $(F, F_{\otimes}, F_*) : \mathcal{C} \rightarrow \mathcal{C}'$ between symmetric monoidal categories \mathcal{C} and \mathcal{C}' is a monoidal functor (F, F_{\otimes}, F_*) such that

$$\begin{array}{ccc}
FA \otimes FB & \xrightarrow{F_{\otimes, A, B}} & F(A \otimes B) \\
\sigma'_{FA, FB} \downarrow & & \downarrow F\sigma_{A, B} \\
FB \otimes FA & \xrightarrow{F_{\otimes, B, A}} & F(B \otimes A)
\end{array}$$

commutes for all $A, B \in \mathcal{C}$.

Definition 1.8 (Monoidal natural transformation). A *monoidal natural transformation* $\theta : F \Rightarrow G$ between two monoidal functors F and G is a natural transformation $\theta : F \Rightarrow G$ such that the triangle

$$\begin{array}{ccc}
FI & \xrightarrow{\theta_I} & GI \\
& \swarrow F_* & \searrow G_* \\
& & I'
\end{array} \tag{MNT1}$$

and square

$$\begin{array}{ccc}
FA \otimes FB & \xrightarrow{\theta_A \otimes \theta_B} & GA \otimes GB \\
F_{\otimes, A, B} \downarrow & & \downarrow G_{\otimes, A, B} \\
F(A \otimes B) & \xrightarrow{\theta_{A \otimes B}} & G(A \otimes B)
\end{array} \tag{MNT2}$$

commute for all objects A, B .

Example 1.9. $\mathbf{FVect}_{\mathbb{R}}$ and $\mathbf{Mat}(\mathbb{R})$ are equivalent via strong monoidal functors. It is a corollary of the Mac Lane Coherence theorems that any monoidal category can be

‘strictified’—that is, for any monoidal category there exists a strict monoidal category equivalent to it via strong monoidal functors—and that a symmetric monoidal category can be strictified into a strict symmetric monoidal category. For more details see [14, Theorem XI.3.1].

1.2 Graphical calculi

One of the draws of expressing concepts as symmetric monoidal categories is that the structure of these categories naturally lends itself to being expressed pictorially. These pictures, known as *string diagrams*, represent the morphisms of a monoidal category, and have the benefit of hiding certain structural equalities and making use of our topological intuitions to suggest other important equalities. The aim of this section is merely to give the reader a basic working understanding of how to read and draw these diagrams; we leave the precise definition of a string diagram and proofs of their expressiveness to the survey [20] of Selinger.

String diagrams are drawn in two dimensions with, roughly speaking, one dimension representing the categorical composition and the other representing monoidal composition. We shall take the convention, common but far from universal, that we read composition up the page, leaving horizontal juxtaposition to represent the monoidal product of maps. Under this convention then, a string diagram consists of a graph with edges labelled by objects and vertices labelled by morphisms, which as a whole represents a morphism with domain the monoidal product of the edges at the lower end of the diagram, and codomain the monoidal product of the edges at the top.

The simplest example, consisting of just a single edge, represents the identity map:

$$\text{id}_X = \begin{array}{c} X \\ | \\ X \end{array}$$

More generally, we represent a morphism $f : X \rightarrow Y$ by drawing in sequence up the page an edge labelled by X , ending at a vertex labelled by f , which then gives rise

to an edge labelled by Y :

$$f = \begin{array}{c} Y \\ | \\ \boxed{f} \\ | \\ X \end{array}$$

If $X = A \otimes B \otimes C$, and $Y = D \otimes E$, we could also represent f as:

$$f = \begin{array}{c} D \quad E \\ | \quad | \\ \boxed{f} \\ | \quad | \quad | \\ A \quad B \quad C \end{array}$$

Given maps $f : X \rightarrow Y$ and $g : Y \rightarrow Z$, we represent their composite $g \circ f$ by placing a vertex g on the Y -edge leaving the vertex f :

$$g \circ f = \begin{array}{c} Z \\ | \\ \boxed{g} \\ | \\ \boxed{f} \\ | \\ X \end{array}$$

If the types of the maps are known, we lose no information if we omit the labels of edges that are connected to the vertices of the maps, as we have done for edge representing Y in the above diagram. For the sake of cleanness and readability, we shall most often just label the ‘input’ and ‘output’ edges at the top and bottom of the diagram.

The monoidal product of two maps is given by their horizontal juxtaposition, with juxtaposition on the right representing monoidal product on the right, and on the left representing left monoidal product. As an example, given morphisms $f : X \rightarrow Y$ and $g : Z \rightarrow W$, we write their product $f \otimes g : X \otimes Z \rightarrow Y \otimes W$ as:

$$f \otimes g = \begin{array}{c} Y \quad W \\ | \quad | \\ \boxed{f} \quad \boxed{g} \\ | \quad | \\ X \quad Z \end{array}$$


The monoidal unit is an object with special properties in the category, and as a result the conventions for representing the unit diagrammatically are a little different: we

don't draw it or its identity map


$$\text{id}_I =$$

This has the advantage of any diagram representing a morphism $f : A \rightarrow B$ also representing the 'equivalent' morphism $f \otimes \text{id}_I : A \otimes I \rightarrow B \otimes I$, among other equivalent morphisms.

To read an arbitrary string diagram, it is often easiest to start at the lower edge and move up the diagram, reading off a morphism for every horizontal cross-section intersecting a vertex. The string diagram then represents the composite of these morphisms in the order that the morphisms were read, applying associators and unitors as needed for the map to be well-defined. For example, reading in this way the diagram



represents the map $(f \otimes g) \circ (h \otimes k)$. Note that it may also be read $(f \circ h) \otimes (g \circ k)$, or even $(f \circ \text{id}_X \circ h) \otimes (g \circ k) \otimes \text{id}_I$ where X is the codomain of f , but in any case all these different algebraic descriptions of the picture represent the same morphism. This is a key feature of string diagrams: many equalities of algebraic representations of morphisms become just the identity of diagrams. Furthermore, we need not be too careful about the precise geometry of the diagrams; the following topologically equivalent diagrams in fact also express equal morphisms:



This holds true in general.

Theorem 1.10 (Coherence of the graphical calculus for monoidal categories). *Two morphisms in a monoidal category are equal with their equality following from the axioms of monoidal categories if and only if their diagrams are equal up to planar deformation.*

Proof. Joyal-Street [11, Theorem 1.2]. □

In a symmetric monoidal category, we usually omit the label for the swap, denoting it instead just by the intersection of two strings:

$$\sigma_{A,B} = \boxed{\sigma_{A,B}} = \begin{array}{c} B \quad A \\ | \quad | \\ \diagdown \quad / \\ | \quad | \\ A \quad B \end{array}$$

We will also later take such an approach for other chosen maps, such as the multiplication and unit of a monoid.

The defining identities of the swap may then be written graphically as

$$\begin{array}{c} A \quad B \\ | \quad | \\ \diagdown \quad / \\ | \quad | \\ A \quad B \end{array} = \begin{array}{c} A \quad B \\ | \quad | \\ | \quad | \\ A \quad B \end{array} \quad (\text{Sym1})$$

and

$$\begin{array}{c} B \quad C \quad A \\ | \quad | \quad | \\ \diagdown \quad / \quad | \\ | \quad | \quad | \\ A \quad B \quad C \end{array} = \begin{array}{c} B \quad C \quad A \\ | \quad | \quad | \\ \diagdown \quad / \quad | \\ | \quad | \quad | \\ A \quad B \quad C \end{array} \quad (\text{Sym2})$$

Including these identity into our collection of allowable transformations of diagrams gives coherence theorem for symmetric monoidal categories.

Theorem 1.11 (Coherence of the graphical calculus for symmetric monoidal categories). *Two morphisms in a symmetric monoidal category are equal with their equality following from the axioms of symmetric monoidal categories if and only if their diagrams are equal up to planar deformation and local applications of the identities Sym1 and Sym2.*

Proof. Joyal-Street [11, Theorem 2.3]. □

Just as two diagrams represent the same morphism in a monoidal category if they agree up to planar isotopy, this theorem may be regarded geometrically as stating that two diagrams represent the same morphism in a monoidal category if they agree up to isotopy in four dimensions.

These two theorems show that the graphical calculi go beyond visualisations of the morphisms, having the ability to provide bona-fide proofs of equalities of morphisms. As a general principle, one which we shall demonstrate in this dissertation, this fact combined the intuitiveness of manipulations and the encoding of certain equalities and structural isomorphisms make the string diagrams better than the conventional algebraic language for understanding monoidal categories.

1.3 Example: the theory of monoids

This section serves to both give examples of the constructions defined in this chapter and, more importantly, give a flavour of the spirit in which we will aim to use monoidal categories to discuss causality.

Recall that a monoid is a set with an associative, unital binary operation. We shall classify these as strong monoidal functors from a category $\mathbf{Th}(\mathbf{Mon})$ into \mathbf{Set} , and hence say that this category $\mathbf{Th}(\mathbf{Mon})$ describes the theory of monoids. The study of this category and its functorial images then gives new and interesting perspectives of the concept of a monoid, its generalisations, and relationships to other mathematical structures. In analogy to this, we will later define causal theories as monoidal categories that can be modelled within other categories through monoidal functors.

Define the category $\mathbf{Th}(\mathbf{Mon})$ as follows: fix some symbol M , and let the objects of $\mathbf{Th}(\mathbf{Mon})$ be any natural number of copies of this symbol. We shall write the objects $M^{\otimes n}$, where $n \in \mathbb{N}$ is the number of copies of M . Then the monoidal product on the objects of $\mathbf{Th}(\mathbf{Mon})$ is just addition of number of copies of M , with $I = M^{\otimes 0}$ the monoidal unit. By definition this is a strict monoidal category, so we need not worry about the structure maps.

In addition to the identity morphism on each object, we also include morphisms $\frown : M \otimes M \rightarrow M$ and $\smile : I \rightarrow M$ and all their composites and products, subject to the relations

$$\begin{array}{c} \text{---} \\ \text{---} \end{array} \begin{array}{c} \text{---} \\ \text{---} \end{array} \begin{array}{c} \text{---} \\ \text{---} \end{array} = \begin{array}{c} \text{---} \\ \text{---} \end{array} \begin{array}{c} \text{---} \\ \text{---} \end{array} \begin{array}{c} \text{---} \\ \text{---} \end{array} \quad (\text{associativity})$$

and

$$\begin{array}{c} \text{---} \\ \text{---} \end{array} \bullet = \text{---} = \begin{array}{c} \text{---} \\ \text{---} \end{array} \bullet \quad (\text{unitality})$$

These equations correspond respectively to the associativity and unitality laws for the monoid.

Now given any monoid $(X, \cdot, 1)$, we can define the strong monoidal functor $F : \mathbf{Th}(\mathbf{Mon}) \rightarrow \mathbf{Set}$ mapping $M^{\otimes n}$ to the n -fold cartesian product X^n of the set X , m to the monoid multiplication function $X \times X \rightarrow X; (x, y) \mapsto x \cdot y$, and e to the function $\{*\} \rightarrow X; * \mapsto 1$ with image the monoid unit. This is well-defined as the relations obeyed by m and e are precisely those required to ensure the monoid operation \cdot is associative and unital. Furthermore, taking the canonical isomorphisms $X^m \times X^n \xrightarrow{\sim} X^{m+n}$ given by the universal property of products, we see that F is a strong monoidal functor.

Conversely, given any strong monoidal functor $(F, F_{\otimes}, F_*) : \mathbf{Th}(\mathbf{Mon}) \rightarrow \mathbf{Set}$, it is straightforward to show, using the naturality of F_{\otimes} and the diagrams obeyed by the definition of a strong monoidal functor, that the triple $(FM, Fm \circ F_{\otimes, M, M}, Fe \circ F_*(*))$ is a well-defined monoid. From here it also can be shown that these two constructions are inverses up to isomorphism, and so we have bijections

$$\left\{ \begin{array}{c} \text{isomorphism classes of} \\ \text{monoids} \end{array} \right\} \longleftrightarrow \left\{ \begin{array}{c} \text{isomorphism classes of} \\ \text{strong monoidal functors} \\ \mathbf{Th}(\mathbf{Mon}) \rightarrow \mathbf{Set} \end{array} \right\}.$$

This shows that the strong monoidal functors from $\mathbf{Th}(\mathbf{Mon})$ to \mathbf{Set} classify all monoids.

In fact, the category $\mathbf{Th}(\mathbf{Mon})$ classifies not only monoids themselves, but also the maps between them. Indeed, given monoids X, X' and corresponding strong monoidal functors F, F' , we also have a bijection

$$\left\{ \begin{array}{c} \text{monoid homomorphisms} \\ X \rightarrow X' \end{array} \right\} \longleftrightarrow \left\{ \begin{array}{c} \text{monoidal natural transformations} \\ F \Rightarrow F' \end{array} \right\}.$$

This bijection sends a monoid homomorphism $\varphi : X \rightarrow X'$ to the monoidal natural transformation defined on $M \in \mathbf{Th}(\mathbf{Mon})$ by $\varphi : FM \rightarrow F'M$. The requirement that monoid homomorphisms preserve the identity corresponds to the triangle MNT1 that monoidal natural transformations must obey, with the requirement that monoid homomorphisms preserve the monoid multiplication corresponds to the square MNT2. It is further possible to show that these bijections respect composition of monoid homomorphisms and monoidal natural transformations. This shows that the category of monoids is equivalent to the category of strong monoidal functors from $\mathbf{Th}(\mathbf{Mon})$ to \mathbf{Set} . It is in this strong sense that $\mathbf{Th}(\mathbf{Mon})$ classifies monoids, and for this reason we call this category the theory of monoids.

One advantage of this perspective is that we may now talk of monoid objects in other monoidal categories, which are often interesting structures in their own right. This often gives insight into the relationships between known mathematical structures. For example, the category of monoid objects in the monoidal category $(\mathbf{Ab}, \otimes, \mathbb{Z})$ of abelian groups with tensor product as monoidal product and \mathbb{Z} as the monoidal unit can be shown to be precisely the category of rings.

We will use this idea of defining generalised monoid-like objects in categories other than \mathbf{Set} in pursuing a categorical definition of a causal theory. In particular, we will be interested in commutative comonoid objects.

Definition 1.12 (Commutative comonoid). As for in defining $\mathbf{Th}(\mathbf{Mon})$, fix a symbol M , and define the symmetric monoidal category $\mathbf{Th}(\mathbf{CComon})$ to be the symmetric monoidal category with objects tensor powers of M and morphisms generated by the swaps and the maps $\Upsilon : M \rightarrow M \otimes M$, $\uparrow : M \rightarrow I$, subject to the relations

$$\begin{array}{c}
 \begin{array}{ccc}
 \begin{array}{c} \diagup \quad \diagdown \\ \diagdown \quad \diagup \\ \diagup \quad \diagdown \\ \diagdown \quad \diagup \\ \uparrow \end{array} & = & \begin{array}{c} \diagup \quad \diagdown \\ \diagdown \quad \diagup \\ \diagup \quad \diagdown \\ \diagdown \quad \diagup \\ \uparrow \end{array} \\
 \end{array} & & \text{(coassociativity)}
 \end{array}$$

$$\begin{array}{c}
 \begin{array}{ccc}
 \begin{array}{c} \diagup \quad \diagdown \\ \diagdown \quad \diagup \\ \cdot \\ \uparrow \end{array} & = & \begin{array}{c} | \\ | \\ | \\ \uparrow \end{array} & = & \begin{array}{c} \cdot \\ \diagup \quad \diagdown \\ \diagdown \quad \diagup \\ \uparrow \end{array} \\
 \end{array} & & \text{(counitality)}
 \end{array}$$

and

$$\begin{array}{c}
 \begin{array}{ccc}
 \begin{array}{c} \diagup \quad \diagdown \\ \diagdown \quad \diagup \\ \diagup \quad \diagdown \\ \diagdown \quad \diagup \\ \uparrow \end{array} & = & \begin{array}{c} \diagup \quad \diagdown \\ \diagdown \quad \diagup \\ \uparrow \end{array} \\
 \end{array} & & \text{(commutativity)}
 \end{array}$$

A *commutative comonoid* in a symmetric monoidal category \mathcal{C} is a strong symmetric monoidal functor $\mathbf{Th}(\mathbf{CComon}) \rightarrow \mathcal{C}$. Abusing our terminology slightly, we will often just say that the image of M under this functor is a commutative comonoid.

Chapter 2

Categorical Probability Theory

Probability theory concerns itself with *random variables*: properties of a system that may take one of a number of possible outcomes, together with a likelihood for each possible observation regarding the property. We call the property itself the *variable*, and together the likelihoods for the observations form a *probability assignment* for the variable. As we will mainly concern ourselves with relationships between random variables, we have particular interest in rules that specify a probability assignment on one variable given a probability assignment on another—this can be seen as the latter variable having some causal influence on the former.

In this chapter we will develop the standard tools to talk about all these things, but with emphasis on a categorical perspective. These categorical ideas originate with Lawvere [13], and were extended by Girly in [9]. We caution that the terminology we have used for the basic concepts in probability is slightly nonstandard, but predominantly follows that of Pearl [19] and the Bayesian networks community. Although it will not affect the mathematics, we will implicitly take a frequentist view of probability to complement our physical interpretation of causality.

2.1 The category of measurable spaces

The idea of a variable is captured by measurable spaces. These consist of a set X , thought of as the set of ‘outcomes’ of the variable, and a collection Σ of subsets of X obeying certain closure properties, which represent possible observations about X and which we call the measurable sets of X . We then talk of probability assignments on these measurable spaces via a function $P : \Sigma \rightarrow [0, 1]$ satisfying some consistency properties. While the collection of measurable sets is often taken to be the power set $\mathcal{P}(X)$ when X is finite, for larger sets some restrictions are usually necessary if one wants to assign interesting collections of probabilities to the space.

Given a measurable set A , we think of the number $P(A)$ as the chance that the outcome of the random variable with outcomes represented by X will lie in the subset A . As an example, the process of rolling a six-sided die can be described by the measurable space with set of outcomes $X = \{1, 2, 3, 4, 5, 6\}$ and measurable subsets $\Sigma = \mathcal{P}(X)$ all subsets of X . The statement that the die is fair is then the statement that the probability associated to any subset $A \subseteq X$ is $P(A) = \frac{1}{6}|A|$.

We formalise this in the following standard way; more details can be found in [3] or [21], or indeed any introductory text to probability theory.

Definitions 2.1 (σ -algebra, measurable space). Given a set X , a σ -algebra Σ on X is a set Σ of subsets of X that contains the empty set and is closed under both countable union and complementation in X . We call a pair (X, Σ) consisting of a set X and a σ -algebra Σ on X a *measurable space*.

On occasion we will just write X for the measurable space (X, Σ) , leaving the σ -algebra implicit. In these cases we will write Σ_X to mean the σ -algebra on X .

Example 2.2 (Discrete and indiscrete measurable spaces). Let X be a set. The power set $\mathcal{P}(X)$ of X forms a σ -algebra, and we call $(X, \mathcal{P}(X))$ a *discrete* measurable space. At the other extreme, distinct whenever X has more than one element, is the σ -algebra $\{\emptyset, X\}$. In this case we call $(X, \{\emptyset, X\})$ an *indiscrete* measurable space.

Even beyond the two of the above example, it is not hard to find σ -algebras: we may construct one from any collection of subsets. Indeed, we say that the σ -algebra $\Sigma(\mathcal{G})$ generated by a collection $\mathcal{G} = \{G_i\}_{i \in I}$ of subsets of a set X is the intersection of all σ -algebras on X containing \mathcal{G} . An explicit construction can be given by taking all countable intersections of the sets in \mathcal{G} and their complements, and then taking all countable unions of the resulting sets. We say that a measurable space is *countably generated* if there exists a countable generating set for it.

Example 2.3 (Borel measurable spaces). Many frequently used examples of measurable spaces come from topological spaces. The *Borel σ -algebra* \mathcal{B}_X of a topological space X is the σ -algebra generated by the collection of open subsets of the space.

Example 2.4 (Product measurable spaces). Given measurable spaces (X, Σ_X) , (Y, Σ_Y) , we write $\Sigma_X \otimes \Sigma_Y$ for the σ -algebra on $X \times Y$ generated by the collection subsets $\{A \times B \subseteq X \times Y \mid A \in \Sigma_X, B \in \Sigma_Y\}$. We call this the *product σ -algebra* of Σ_X and Σ_Y , and call the resulting measurable space $(X \times Y, \Sigma_X \otimes \Sigma_Y)$ the *product measurable*

space of (X, Σ_X) and (Y, Σ_Y) . Note that as $(A \times B) \cap (A' \times B') = (A \cap A') \times (B \cap B')$ and $(A \times B)^c = (A \times B^c) \cup (A^c \times B)$, we may write

$$\Sigma_X \otimes \Sigma_Y = \left\{ \bigcup_{i \in I} (A_i \times B_i) \mid A_i \in \Sigma_X, B_i \in \Sigma_Y \right\}.$$

The product measurable space is in fact a categorical product in the category of measurable spaces. To understand this, we first must specify the notion of morphism corresponding to measurable spaces. Just as continuous functions reflect the open sets of a topology, the important notion of map for measurable sets is that of functions that reflect measurable sets.

Definition 2.5 (Measurable function). A function $f : X \rightarrow Y$ between measure spaces (X, Σ_X) and (Y, Σ_Y) is called *measurable* if for each $A \in \Sigma_Y$, $f^{-1}(A) \in \Sigma_X$.

We write **Meas** for the category of measurable spaces and measurable functions. It is easily checked that this indeed forms a category with composition simply composition of functions.

It is also not difficult to check that the product measurable space $(X \times Y, \Sigma_X \otimes \Sigma_Y)$ is the product of the measurable spaces (X, Σ_X) and (Y, Σ_Y) in this category. As the projection maps $\pi_X : X \times Y \rightarrow X$ and $\pi_Y : X \times Y \rightarrow Y$ of the set product are measurable maps, it is enough to show that for any measurable space (Z, Σ_Z) and pair of measurable functions $f : (Z, \Sigma_Z) \rightarrow (X, \Sigma_X)$, and $g : (Z, \Sigma_Z) \rightarrow (Y, \Sigma_Y)$ the unique function $\langle f, g \rangle : Z \rightarrow X \times Y$ given by the product in **Set** is a measurable function. Since for all countable collections $\{A_i \times B_i\}_{i \in I}$ of subsets of $X \times Y$ we have

$$\langle f, g \rangle^{-1} \left(\bigcup_{i \in I} (A_i \times B_i) \right) = \bigcup_{i \in I} \langle f, g \rangle^{-1} (A_i \times B_i) = \bigcup_{i \in I} (f^{-1}(A_i) \cap g^{-1}(B_i)),$$

this is indeed true.

Note also that any one point set $1 = \{*\}$ with its only possible σ -algebra $\{\emptyset, 1\}$ is a terminal object in **Meas**. We thus may immediately view **Meas** as a symmetric monoidal category, with the symmetric monoidal structure given by the fact that **Meas** has finite products. The swaps $\sigma_{X,Y} : X \times Y \rightarrow Y \times X; (x, y) \mapsto (y, x)$ of **Meas** are the same as those of the symmetric monoidal category **Set**. We shall by default consider **Meas** as a symmetric monoidal category in this way.

Similarly, we may also show that the full subcategories **FinMeas** and **CGMeas** with objects finite measurable spaces and countably generated measurable spaces respectively are also a symmetric monoidal category with monoidal product the categorical product.

2.2 Measures and integration

The reason we deal with measurable spaces is that these form the basic structure required for an object to carry some idea of a probability distribution. More precisely, we deal with measurable spaces because they can be endowed with probability measures.

Definitions 2.6 (Measure, measure space). Given a measurable space (X, Σ) , a *measure* μ on (X, Σ) is a function $\mu : \Sigma \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ such that:

- (i) the empty set \emptyset has measure $\mu(\emptyset) = 0$; and
- (ii) if $\{A_i\}_{i \in I}$ is a countable collection of disjoint measurable sets then $\mu(\cup_{i \in I} A_i) = \sum_{i \in I} \mu(A_i)$.

Any such triple (X, Σ, μ) is then known as a *measure space*. When $\mu(X) = 1$, we further call μ a *probability measure*, and (X, Σ, μ) a *probability space*.

We will have to pay close attention to the properties of the collections, in fact σ -ideals, of sets of measure zero of probability spaces in the following. These represent possible observations of our random variable that nonetheless are ‘never’ observed, giving us very little information about their causal consequences. Very often we will pronounce functions equal ‘almost everywhere’ if they agree but for a set of a measure zero. More generally, we say a property with respect to a measure space is true *almost everywhere* or for *almost all* values if it holds except on a set of measure zero. We also say that a measure space is of *full support* if its only subset of measure zero is the empty set \emptyset . Such spaces are necessarily countable measure spaces.

Example 2.7 (Finite and countable measurable spaces). We shall say that a measurable space (X, Σ) is a *finite measurable space* if Σ is a finite set. In this case there exists a finite generating set $\{A_1, \dots, A_n\}$ for Σ consisting of pairwise disjoint subsets of X , and measures μ on (X, Σ) are in one-to-one correspondence with to functions $m : \{A_1, \dots, A_n\} \rightarrow \mathbb{R}_{\geq 0}$, with $\mu(A) = \sum_{A_i \subseteq A} m(A_i)$ for all measurable subsets A of X . Measures may thus also be thought of as vectors with non-negative entries in \mathbb{R}^n , with probability measures those vectors whose entries also sum to 1. We may similarly define *countable measurable spaces*, and note that measures on these spaces are in one-to-one correspondence with functions $m : \mathbb{N} \rightarrow \mathbb{R}_{\geq 0}$.

Writing n for some chosen set with $n \in \mathbb{N}$ elements, note that this suggests each finite measurable space is in some sense ‘isomorphic’ to $n = (n, \mathcal{P}(n))$ for some n . Although this is not true in **Meas**, we will work towards constructing a category in which this is true.

We give two more useful examples of measures.

Example 2.8 (Borel measures, Lebesgue measure). A *Borel measure* is a measure on a Borel measurable space. An important collection of examples of these are the Lebesgue measures on $(\mathbb{R}^n, \mathcal{B}_{\mathbb{R}^n})$. These may be characterised as the unique Borel measure on \mathbb{R}^n such that the measure of each closed n -dimensional cube is given by its n -dimensional volume. See any basic text on measure theory, such as [21, Chapter 1], for more details.

When speaking of \mathbb{R} as a measure space, we will mean \mathbb{R} with its Borel σ -algebra and Lebesgue measure. In particular, when referring to a real-valued measurable function, we shall take the codomain as having this structure.

Example 2.9 (Product measures). Given measure spaces (X, Σ_X, μ) and (Y, Σ_Y, ν) , we may define the *product measure* $\mu \times \nu$ on the product measurable space $(X \times Y, \Sigma_X \otimes \Sigma_Y)$ as the unique measure on this space such that for all $A \in \Sigma_X$ and $B \in \Sigma_Y$,

$$\mu \times \nu(A \times B) = \mu(A)\nu(B).$$

A proof of the existence and uniqueness of such a measure can be found in [21, Theorem 6.1.5].

One way in which measures interact with measurable functions is that measures may be ‘pushed forward’ from the domain to the codomain of a measurable map.

Definition 2.10 (Push-forward measure). Let (X, Σ_X, μ) measure space, (Y, Σ_Y) measurable space, and $f : X \rightarrow Y$ be a measurable function. We then define the *push-forward measure* μ_f of μ along f to be the map $\Sigma_Y \rightarrow \mathbb{R}$ given by

$$\mu_f(B) = \mu(f^{-1}(B)).$$

Note that $\mu_f(Y) = \mu(f^{-1}(Y)) = \mu(X)$, so the push-forward of a probability measure is again a probability measure.

As causality concerns the relationships between random variables, we shall be particularly interested in measures on product spaces, so-called *joint measures*. An important example of a push-forward measure is that of the *marginals* of a joint measure. These are the push-forward measures of a joint measure along the projections of the product space: given a joint measure space $(X \times Y, \Sigma_X \otimes \Sigma_Y, \mu)$ with projections $\pi_X : X \times Y \rightarrow X$ and $\pi_Y : X \times Y \rightarrow Y$, we define the marginal μ_X of μ on X to be the push forward measure of μ along π_X , and similarly for μ_Y . We also say that we have

marginalised over Y when constructing the marginal μ_X from the measure μ . Note that the marginals of a joint probability measure are again probability measures.

Observe that for each point x in its domain, a measurable function $f : (X, \Sigma_X) \rightarrow (Y, \Sigma_Y)$ induces a ‘point measure’

$$\delta_{f,x}(B) = \begin{cases} 1 & \text{if } f(x) \in B, \\ 0 & \text{if } f(x) \notin B, \end{cases}$$

on its codomain (Y, Σ_Y) . From this point of view, the push-forward measure of some measure μ on (X, Σ_X) along f can be seen as taking the ‘ μ -weighted average’ or ‘expected value’ of these induced point measures on (Y, Σ_Y) . More precisely, the push-forward measure may be defined as the integral of these point measures with respect to μ .¹

For the sake of completeness, we quickly review the definition of the integral for bounded real-valued measurable functions nonzero on a set of finite measure, but the reader is referred to Ash [3, §1.5] or Stein and Shakarchi [21, Chapter 2] for full detail.

We first define the integral of simple functions. Let (X, Σ, μ) be a measure space, and let A be subset of X . We write $\chi_A : X \rightarrow \mathbb{R}$ for the characteristic function

$$\chi(x) = \begin{cases} 1 & \text{if } x \in A; \\ 0 & \text{if } x \notin A, \end{cases}$$

and call a weighted sum $\varphi = \sum_{k=1}^N c_k \chi_{A_k}$ of characteristic functions of measurable sets a *simple function*. The integral $\int_A \varphi d\mu$ of a simple function φ over the measurable set A with respect to μ is defined to be

$$\int_A \varphi d\mu = \sum_{k=1}^N c_k \mu(A_k \cap A)$$

when this sum is finite. Note that this implies that the integral over X of the characteristic function of a measurable set A is just $\mu(A)$.

Let now f be a bounded real-valued measurable function such that the set $\{x \in X \mid f(x) \neq 0\}$ is of finite measure. It can be shown there then exists a uniformly bounded sequence $\{\varphi_n\}_{n \in \mathbb{N}}$ of simple functions supported on the support of f and

¹A complementary perspective views the integral in terms of push-forwards, but only once we have define the standard notion of multiplying functions with measures to produce a new measure. Indeed, given a bounded real-valued measurable function f and a measure μ on a measurable space (X, Σ) , this new measure $f\mu$ is equal to $\int_A f d\mu$ on each $A \in \Sigma$, and this allows us to see the integral $\int_A f d\mu$ as the value, on the set $*$, of the push-forward measure of the measure $\chi_A f\mu$ along the unique map $X \rightarrow *$ to the terminal object.

converging to f for almost all x . Using this sequence, we define the integral $\int_A f d\mu$ of f over A with respect to μ to be

$$\int_A f d\mu = \lim_{n \rightarrow \infty} \int_A \varphi_n d\mu.$$

By our assumptions, this limit always exists, is finite, and is independent of the sequence $\{\varphi_n\}_{n \in \mathbb{N}}$. Where we do not write the domain of integration A , we mean that the integral is taken over the entire domain X of f .

We will not discuss the technicalities of the integral further, but instead note that in the case of Lebesgue measure the notion of integration agrees with that of Riemann integration, and for finite measure spaces it can be viewed as analogous to matrix multiplication—this will be explained fully in the following section. Our examples will be limited to these cases.

More generally, this idea of averaging measures will play a crucial role in how we reason about consequences of causal relationships. As an illustration, suppose that we have measurable spaces C and R , representing say cloud cover and rain on a given day respectively, and for each value of cloud cover—that is, each measurable set in C —we are given the probability of rain. We will assume this forms a real-valued measurable function f on C . If we are further given a measure μ on C representing how cloudy a day is likely to be, we can ‘average’ over this measure to give a probability of rain on that day. This averaging process is given by the integral of f with respect to μ .

Implicitly here we are talking about conditional probabilities—for each outcome of the space C we get a measure on R . This idea will form our main idea of map between measurable spaces.

2.3 The category of stochastic maps

Measurable functions describe a deterministic relationship between two variables: if one fixes an outcome of the domain variable, a measurable function specifies a unique corresponding outcome for the codomain. When describing a more stochastic world, such as that given by a Markov chain, such certainty is often out of reach. In these cases stochastic maps—variously also called stochastic kernels, Markov kernels, conditional probabilities, or probabilistic mappings—may often be useful instead. These are more general, mapping outcomes of the domain to probability measures on, instead of points of, the codomain.

Definition 2.11 (Stochastic map). Let (X, Σ_X) and (Y, Σ_Y) be measurable spaces. A *stochastic map* $k : (X, \Sigma_X) \rightarrow (Y, \Sigma_Y)$ is a function

$$k(x, B) : X \times \Sigma_Y \longrightarrow [0, 1]$$

such that

- (i) for each $x \in X$ the function $k_x := k(x, -) : \Sigma_Y \rightarrow [0, 1]$ is a probability measure on Y ; and
- (ii) for each measurable set $B \subseteq Y$ the function $k_B := k(-, B) : X \rightarrow [0, 1]$ is measurable.

The *composite of stochastic maps*

$$\ell \circ k : (X, \Sigma_X) \xrightarrow{k} (Y, \Sigma_Y) \xrightarrow{\ell} (Z, \Sigma_Z)$$

is defined by the integral

$$\ell \circ k(x, C) = \int_Y \ell(-, C) dk_x,$$

where $x \in X$ and $C \in \Sigma_Z$. That this is a well-defined stochastic map follows immediately from the basic properties of the integral.

Note that these definitions are those suggested by our discussion at the close of the previous section: put more succinctly, a stochastic map is a measure-valued function (subject to a measurability requirement), and the composite $\ell \circ k(x, C)$ of stochastic maps ℓ and k is given by integrating the measures ℓ_y on the codomain (Z, Σ_Z) with respect to the measure k_x on the intermediate variable (Y, Σ_Y) .

We give a few examples.

Example 2.12 (Probability measures as stochastic maps). Observe that a stochastic map $k : 1 \rightarrow (X, \Sigma)$ is simply a probability measure on (X, Σ) .

Example 2.13 (Deterministic stochastic maps). In the previous section we discussed how a measurable function $f : (X, \Sigma_X) \rightarrow (Y, \Sigma_Y)$ induces ‘point measures’ on its codomain. We can now interpret these as defining the stochastic map $\delta_f : (X, \Sigma_X) \rightarrow (Y, \Sigma_Y)$ given by

$$\begin{aligned} \delta_f : X \times \Sigma_Y &\longrightarrow [0, 1]; \\ (x, B) &\longmapsto \delta_{f,x}(B) = \begin{cases} 1 & \text{if } f(x) \in B; \\ 0 & \text{if } f(x) \notin B. \end{cases} \end{aligned}$$

We call this the *deterministic stochastic map induced by f* . More generally, we call any stochastic map taking values in only in the set $\{0, 1\}$ a *deterministic stochastic map*.

Observe that given measurable functions $f : (X, \Sigma_X) \rightarrow (Y, \Sigma_Y)$ and $g : (Y, \Sigma_Y) \rightarrow (Z, \Sigma_Z)$, the composite of their induced maps is given by

$$\delta_g \circ \delta_f(x, C) = \int_Y \delta_g(-, C) d\delta_{f,x} = \delta_{f,x}(g^{-1}(C)) = \begin{cases} 1 & \text{if } g \circ f(x) \in C; \\ 0 & \text{if } g \circ f(x) \notin C, \end{cases}$$

where $x \in X$ and $C \in \Sigma_Z$. Thus $\delta_g \circ \delta_f = \delta_{g \circ f}$.

More generally, for a stochastic map k and measurable function f of the types required for composition to be well-defined, we have $k \circ \delta_f(x, B) = k(f(x), B)$, and $\delta_f \circ k(x, B) = k(x, f^{-1}(B))$.

Example 2.14 (Stochastic matrices). Let X and Y be finite measurable spaces of cardinality n and m respectively. Note that if $x, x' \in X$ are such that x lies in a measurable set A if and only if x' does, then for any stochastic map $k : X \rightarrow Y$ the measurability of k_B for each $B \in \Sigma_Y$ implies the measures k_x and $k_{x'}$ must be equal. Thus, with reference to Example 2.7, we may assume without loss of generality X and Y are discrete. Then, observing that all maps with discrete domain are measurable and recalling that probability distributions on a finite discrete measurable space m may be considered as vectors in \mathbb{R}^m with non-negative entries that sum to one, we see that stochastic maps $k : X \rightarrow Y$ may be considered as $m \times n$ matrices K with non-negative entries and columns summing to one. Indeed, the correspondence is given by having the yx th entry $K_{y,x}$ of K equal to $k(x, \{y\})$ for all $x \in X$ and $y \in Y$. We call such matrices—matrices with entries in $[0, 1]$ and columns summing to 1—*stochastic matrices*.

Let also Z be a discrete finite measurable space, and let $\ell : Y \rightarrow Z$ be a stochastic map, with corresponding stochastic matrix L . Then for all $x \in X$ and $z \in Z$, we have

$$\ell \circ k(x, \{z\}) = \int_Y \ell(-, \{z\}) dk_x = \sum_{y \in Y} \ell(y, \{z\}) k(x, \{y\}),$$

and writing this in matrix notation then gives

$$\ell \circ k(x, \{z\}) = \sum_{y \in Y} L_{z,y} K_{y,x} = (LK)_{z,x}.$$

Thus our representation of finite stochastic maps as stochastic matrices respects composition. This hints at an equivalence of categories.

We are now in a position to define the main category of interest: let the category of stochastic maps, denoted **Stoch**, be the category with objects measurable spaces and morphisms stochastic maps. It is straightforward to show this is a well-defined category. In particular, the associativity of the composition rule follows directly from the monotone convergence theorem [9, Theorem 1], and for each object (X, Σ) of **Stoch** the delta function $\delta : (X, \Sigma) \rightarrow (X, \Sigma)$ defined by

$$\delta(x, A) = \begin{cases} 1 & \text{if } x \in A; \\ 0 & \text{if } x \notin A, \end{cases}$$

—that is, the deterministic stochastic map induced by the identity function on X —is the identity map.

Viewed with this new category, Example 2.13 defines a functor $\delta : \mathbf{Meas} \rightarrow \mathbf{Stoch}$. In fact, we may further endow **Stoch** with a symmetric monoidal structure such that this is a symmetric monoidal functor. For this we take the product of two objects to be their product measurable space, and the product

$$k \otimes \ell : (X \times Z, \Sigma_X \otimes \Sigma_Z) \rightarrow (Y \times W, \Sigma_Y \otimes \Sigma_W)$$

of two stochastic maps $k : (X, \Sigma_X) \rightarrow (Y, \Sigma_Y)$ and $\ell : (Z, \Sigma_Z) \rightarrow (W, \Sigma_W)$ to be the unique stochastic map extending

$$k \otimes \ell((x, z), B \times D) = k(x, B)\ell(z, D),$$

where $x \in X$, $z \in Z$, $B \in \Sigma_Y$ and $D \in \Sigma_W$. This assigns to each pair (x, z) the product measure of k_x and ℓ_z on $Y \times W$, and indeed results in a well-defined functor $\otimes : \mathbf{Stoch} \times \mathbf{Stoch} \rightarrow \mathbf{Stoch}$. Using as structural maps the induced deterministic stochastic maps of the corresponding structural maps in **Meas** then gives **Stoch** the promised symmetric monoidal structure.

Remark 2.15. Observe that any indiscrete σ -algebra $(*, \{\emptyset, *\})$ is a terminal object in **Stoch**: from any other measurable space (X, Σ) there only exists the map

$$t(x, B) = \begin{cases} 1 & \text{if } B = *; \\ 0 & \text{if } B = \emptyset. \end{cases}$$

Example 2.12 thus shows that the points of an object of **Stoch** are precisely the probability measures on that space.

While **Stoch** has a straightforward definition and interpretation, the generality of the concept of a σ -algebra means that **Stoch** admits a few pathological examples that

indicate it includes more than what we want to capture. For this reason, and for the clarity that simpler cases can bring, we will mostly work with two full subcategories of **Stoch**. The first is **FinStoch**, the category of finite measurable spaces and stochastic maps. Building on Example 2.14, and as promised in Example 2.7, this is monoidally equivalent to the skeletal symmetric monoidal category **SMat** with objects natural numbers and morphisms stochastic matrices. As categories of vector spaces are well studied, this characterisation gives much insight into the structure of **FinStoch**.

The main disadvantage of **FinStoch** is that many random variables are not finite. One category admitting infinite measure spaces—and used by Giry [9], Panangaden [15], and Doberkat [6], among others—is the category of standard Borel spaces,² which can be skeletalised as the countable measurable spaces and the unit interval with its Borel σ -algebra. We will favour the less frequently used but slightly more general category **CGStoch**, the full subcategory of **Stoch** obtained by restricting the objects to the countably generated measurable spaces. This setting is general enough to handle almost all examples of probability spaces that arise in applications, but has a few nice properties that **Stoch** does not. In the next section we see one of them: the deterministic stochastic maps here are precisely those that arise from measurable functions.

2.4 Deterministic stochastic maps

Recall that the deterministic stochastic maps are those that take only the values 0 and 1. These will play a crucial role in maps between collections of causally related random variables. The key reason for this is that these maps show much more respect for the structure of the measurable spaces than general stochastic maps. For example, for a stochastic map to be an isomorphism in **Stoch**, it must be deterministic.

Proposition 2.16. *Let $k : (X, \Sigma_X) \rightarrow (Y, \Sigma_Y)$ be an isomorphism in **Stoch**. Then k is deterministic.*

Proof. Our argument rests on the fact that if $f : X \rightarrow [0, 1]$ is a measurable function on a probability space (X, Σ, μ) such that $\int f d\mu = 1$, then $\mu(f^{-1}\{1\}) = 1$.

Write h for the inverse stochastic map to k , and fix $B \in \Sigma_Y$. We begin by defining $A = k_B^{-1}\{1\}$, where we remind the reader that k_B is the measurable function

²A measurable space is a standard Borel space if it is the Borel measurable space of some Polish space. A topological space is a Polish space if it is the underlying topological space of some complete separable metric space. This category then has objects standard Borel spaces and morphisms stochastic maps between them.

$k(-, B) : X \rightarrow [0, 1]$. Note that we then have $B \subseteq h_A^{-1}\{1\}$, since for any $y \in B$ that h is the inverse to k gives $\int k(-, B) dh_y = 1$, so by the above fact $h_y(A) = 1$, and hence $y \in h_A^{-1}\{1\}$.

It is enough to show that for any $x \in X$, $k(x, B) = 0$ or 1 . If $x \in A$ we are done: by definition then $k(x, B) = 1$. Suppose otherwise. Then, again as h and k are inverses, $\int h(-, A) dk_x = 0$. But

$$\int h(-, A) dk_x \geq k_x(h_A^{-1}\{1\}) \geq k_x(B).$$

Thus $k(x, B) = 0$, as required. \square

In the previous section, we showed that every measurable function induces a deterministic stochastic map. One of the reasons that we prefer to work with countably generated measurable spaces is that in **CGStoch** the converse is also true.

Proposition 2.17. *Let (Y, Σ_Y) be a measurable space with Σ_Y countably generated. Then a stochastic map $k : X \rightarrow Y$ is deterministic if and only if there exists a measurable function $f : X \rightarrow Y$ with $k = \delta_f$.*

Proof. A proof can be found in [5, Proposition 2.1], but we outline a version here to demonstrate the use of the countable generating set, and point out that we assume the axiom of choice.

We have seen that measurable functions induce deterministic maps. For the converse, let \mathcal{G} be a countable generating set for Σ_Y . Now for each $x \in X$ let $B_x = \bigcap_{\{B \in \mathcal{G} \mid k(x, B) = 1\}} B$. This is a measurable set as \mathcal{G} is countable, and has k_x -measure 1 as its complement may be written as a countable union of sets of k_x -measure zero. Choosing then for each x some $y \in B_x$, we define f such that $f(x) = y$. It is then easily checked that $k = \delta_f$, and f is measurable as each k_B is. \square

Remark 2.18. On the other hand, one need not look too hard for a deterministic stochastic map that is not induced by a measurable function when dealing with non-countably generated measurable spaces.³ Indeed, take any uncountable set X , and endow it with the σ -algebra generated by the points of X . This means that set is measurable if and only if it or its complement is countable. It is then easily checked

³For fun, we note that if we further add the requirement that every subset of our codomain be measurable, then we *do* need to look quite hard. We say that a cardinal is a measurable cardinal if there exists a countably-additive two-valued measure on its power set such that it has measure 1 and each point has measure 0. If we are looking for such measures, then our set has to be a strongly inaccessible cardinal [22]. These are truly huge; in some models of set theory they're too huge to exist!

that assigning countable sets measure 0 and uncountable measurable sets measure 1 defines a measure. This gives a deterministic stochastic map from the terminal object to X not induced by any measurable function.

Remark 2.19. Note that the measurable function specifying a deterministic stochastic map need not be unique, so we should not view the deterministic stochastic maps as merely the collection of measurable maps lying inside **CGStoch**. As an example of this, consider the one point measurable space $(*, \{\emptyset, *\})$ and any other indiscrete measurable space (X, Σ) . Then all of the $|X|$ functions $f : * \rightarrow X$ are measurable, and all induce the deterministic stochastic map $\delta_f(*, \emptyset) = 0$; $\delta_f(*, X) = 1$. In this way **Stoch** captures the intuition that every indiscrete measurable space is the same.

In particular, non-bijective measurable endofunctions can induce the identity stochastic map, so measurable spaces may be isomorphic in **Stoch** even if they are not isomorphic in **Meas**. This lets **FinStoch** admit the skeletalisation **SMat**, even while the classification of isomorphic objects in **FinMeas** is not nearly so neat.

More abstractly, this shows that although our symmetric monoidal functor $\delta : \mathbf{Meas} \rightarrow \mathbf{Stoch}$ is injective on objects, it is not faithful, and so we can not view **Meas** as a subcategory of **Stoch**.

Although inducing deterministic stochastic maps from measurable functions is in general a many-to-one process, we may always take quotients of our measurable spaces so it becomes one-to-one. We briefly explore this idea in order to further our understanding of deterministic stochastic maps.

Call two outcomes of a measurable space *distinguishable* if there exists a measurable subset containing one but not the other, and *indistinguishable* otherwise. Indistinguishability gives an equivalence relation on the outcomes of a measure space. We may take a quotient by this equivalence relation, and use the quotient map to induce a σ -algebra on the quotient set, defining a set in the quotient to be measurable if its preimage is. In this quotient space all outcomes are distinguishable; we call this an *empirical measurable space*. The quotient map in fact induces an isomorphism in **Stoch**.

We call a deterministic monomorphism in **Stoch** an *embedding* of measurable spaces. These are deterministic stochastic maps induced by injective measurable functions on the empiricisations, and so may be thought of as maps that realise a measurable space as isomorphic to a sub-measurable space of another. We call an epimorphism in **Stoch** a *coarse graining* of measurable spaces. These are deterministic stochastic maps induced by surjective measurable functions on the empiricisations,

and so may be thought of as maps that remove the ability to distinguish between some outcomes of the domain.

The following proposition then gives a precise understanding of deterministic stochastic maps in **CGStoch**.

Proposition 2.20. *In **CGStoch**, every deterministic stochastic map may be factored as a coarse graining followed by an embedding.*

Proof. We may without loss of generality assume spaces are empirical. Then we may treat the deterministic stochastic maps as functions, and we know that each function factors into a surjection followed by an injection. \square

2.5 Aside: the Giry monad

To shed further light on the close relationship between **Meas** and **Stoch**, we mention a few results that first stated in [13], and proved in [9]. The main observation is that **Stoch** forms a relation-like version of **Meas**. More precisely, we observe that just as **Rel** is the Kleisli category for the power set monad on **Set**, **Stoch** is the Kleisli category for Giry monad on **Meas**.

Recall that a monad on a category \mathcal{C} consists of a functor $T : \mathcal{C} \rightarrow \mathcal{C}$ and natural transformations $\eta : 1_{\mathcal{C}} \Rightarrow T$ and $\mu : T^2 \Rightarrow T$ such that for all objects $X \in \mathcal{C}$ the diagrams

$$\begin{array}{ccc}
 T(T(T(X))) & \xrightarrow{T(\mu_X)} & T(T(X)) \\
 \mu_{T(X)} \downarrow & & \downarrow \mu_X \\
 T(T(X)) & \xrightarrow{\mu_X} & T(X)
 \end{array}
 \quad \text{and} \quad
 \begin{array}{ccc}
 T(X) & \xrightarrow{\eta_{T(X)}} & T(T(X)) \\
 T(\eta_X) \downarrow & \searrow & \downarrow \mu_X \\
 T(T(X)) & \xrightarrow{\mu_X} & T(X)
 \end{array}$$

commute. Also recall that the Kleisli \mathcal{C}_T category of such a monad on \mathcal{C} is the category with objects that of \mathcal{C} , for all $X, Y \in \mathcal{C}$ homsets $\text{hom}_{\mathcal{C}}(X, TY)$, and composition of $f^* : X_T \rightarrow Y_T$, $g^* : Y_T \rightarrow Z_T$, defined by $f : X \rightarrow TY$, $g : Y \rightarrow TZ$, given by $g^* \circ_T f^* = (\mu \circ Tg \circ f)^*$.

As mentioned above, it can be checked that the functor mapping a set to its power set can be viewed as a monad on **Set**, and the Kleisli category for this monad is isomorphic to **Rel**. In the case of **Meas** and **Stoch**, we define the functor of the Giry monad $\mathcal{P} : \mathbf{Meas} \rightarrow \mathbf{Meas}$ to be the functor taking a measurable space (X, Σ)

to the set of all probability measures on (X, Σ) with the smallest σ -algebra such that the evaluation maps

$$\begin{aligned} \mathcal{P}X \times \Sigma &\longrightarrow [0, 1]; \\ (\mu, B) &\longmapsto \mu(B). \end{aligned}$$

are measurable.⁴ The associated natural transformations of the monad are that $1 \rightarrow \mathcal{P}$ sending a point to its point measure, and that $\mathcal{P}^2 \rightarrow \mathcal{P}$ sending a measure on the set of measures to its integral. It can then be shown that this forms a well-defined monad, with Kleisli category **Stoch**.

As **Stoch** is the Kleisli category for \mathcal{P} , \mathcal{P} can be factored through **Stoch**, and in fact through the functor $\delta : \mathbf{Meas} \rightarrow \mathbf{Stoch}$. This is done by defining the functor $\varepsilon : \mathbf{Stoch} \rightarrow \mathbf{Meas}$ sending a measurable space X to $\mathcal{P}X$ and a stochastic map $k : X \rightarrow Y$ to the measurable function $\varepsilon k : \mathcal{P}X \rightarrow \mathcal{P}Y$ defined by

$$\begin{aligned} \varepsilon k : \mathcal{P}X &\longrightarrow \mathcal{P}Y; \\ \mu &\longmapsto (B \in \Sigma_Y \mapsto \int_X k(x, B) d\mu). \end{aligned}$$

We then have an adjunction $\delta \dashv \varepsilon$

$$\mathbf{Meas} \begin{array}{c} \xrightarrow{\delta} \\ \xleftarrow{\varepsilon} \end{array} \mathbf{Stoch}$$

with composite $\gamma \circ \delta = \mathcal{P}$.

Finally, note that if X is finite or countably generated then $\mathcal{P}X$ is finite or countably countably generated respectively too, so we may also view **FinStoch** and **CGStoch** as Kleisli categories of monads on **FinMeas** and **CGMeas** respectively.

⁴We earlier saw hints that a stochastic map may be viewed as a measure-valued measurable function. We now see the precise meaning of this statement: a stochastic map is defined by a measurable function $X \rightarrow \mathcal{P}X$.

Chapter 3

Bayesian Networks

In the first chapter we discussed a formalism for representing processes, while in the second we introduced a way to think of these processes as probabilistic. In this short third chapter we now add to this some language for describing selected probabilistic processes as causal.

As in the case of probability, although the intuition for the concept is clear, any attempt to make precise what is meant by causality throws up a number of philosophical questions. We shall not delve into these here, but instead say that we will naively view a causal relationship as an asymmetric one between two variables, in which the varying of one—the cause—necessarily induces variations in the other—the effect. In particular, we think of a causal relationship as implying a physical, objective, mechanism through which this occurs.

The structure we have chosen, Bayesian networks, has roots in graphical models in statistics, and was first proposed as a language for causality by Pearl in [17], with special interest in applications to machine learning and artificial intelligence. Since then Bayesian networks have played a significant role in discussions of causality from both a computational and a philosophical perspective. This chapter in particular relies on expositions by Pearl [19] and Williamson [26].

3.1 Conditionals and independence

Much of the difficulty in the discussion of causality arises from the fact that causal relationships can never be directly observed. We instead must reconstruct such relationships from hints in independencies between random variables. The key point is that if A causes B , then A and B cannot be independent.

Definition 3.1 (Independence). Let (X, Σ_X) and (Y, Σ_Y) be measurable spaces, and let P be a joint probability measure on the product space $(X \times Y, \Sigma_X \otimes \Sigma_Y)$. We say that X and Y are *independent with respect to P* if P is equal to the product measure of P_X and P_Y , and *dependent with respect to P* otherwise.

Independent joint distributions can also be characterised as those that are of the form

$$P(A \times B) = \int_A c(-, B) dP_X,$$

for all $A \in \Sigma_X$ and $B \in \Sigma_Y$, where $c : X \rightarrow Y$ is a stochastic map that factors through the terminal object of **Stoch**. Indeed, to find such a c corresponding to any independent joint distribution, we may just take the stochastic map defined by

$$c(x, B) = P_Y(B)$$

for all $x \in X$. This general idea gives a recipe for deconstructing, or *factorising*, a joint probability measure P into a marginal P_X on one factor and a stochastic map $P_{Y|X} : X \rightarrow Y$ from that factor to the product of the others. We call this stochastic map a conditional for the joint measure.

Definition 3.2 (Conditional). Let (X, Σ_X) and (Y, Σ_Y) be measurable spaces, and let P be a joint probability measure on the product space $(X \times Y, \Sigma_X \otimes \Sigma_Y)$. Then we say that a stochastic map $P_{Y|X} : (X, \Sigma_X) \rightarrow (Y, \Sigma_Y)$ is a *conditional for P with respect to X* if for all $A \in \Sigma_X$ and $B \in \Sigma_Y$ we have

$$P(A \times B) = \int_A P_{Y|X}(-, B) dP_X.$$

Note that the above integral consequently defines a measure on $X \times Y$ equal to P .

Considering the marginals as stochastic maps $1 \rightarrow (X, \Sigma_X)$, this also implies that

$$P_Y = P_{Y|X} \circ P_X.$$

This says that $P_{Y|X}$ is a stochastic map from X to Y that maps the marginal P_X to the marginal P_Y . As there are many joint measures with marginals P_X and P_Y , however, this is not a sufficient condition for $P_{Y|X}$ to be the conditional for P with respect to X .

While this is so, given a joint probability measure with marginals again probability measures, under mild constraints it is always true that there exists a conditional for it, and that this conditional is ‘almost’ unique. This is made precise by the following

proposition. Recall that given a measurable space (X, Σ) , we call a measure μ on X *perfect* if for any measurable function $f : X \rightarrow \mathbb{R}$ there exists a Borel measurable set $E \subseteq f(X)$ such that $\mu(f^{-1}(E)) = \mu(X)$. This proposition represents another reason why we will occasionally restrict our attention to **CGStoch**.

Proposition 3.3 (Existence of regular conditionals). *Let (X, Σ_X) and (Y, Σ_Y) be countably generated measurable spaces, and let μ be a measure on the product space such that the marginal μ_X is perfect. Then there exists a stochastic map $k : (X, \Sigma_X) \rightarrow (Y, \Sigma_Y)$ such that for all $A \in \Sigma_X$ and $B \in \Sigma_Y$ we have*

$$\mu(A \times B) = \int_A k(-, B) d\mu_X.$$

Furthermore, this stochastic map is unique in the sense that if k' is another stochastic map with these properties, then k and k' are equal μ_X almost everywhere.

Proof. Existence is proved in Faden [7, Theorem 6]. Uniqueness is not difficult to show, but can be found in Vakhania, Tarieladze [23, Proposition 3.2]. \square

The existence of conditionals gives rise to a more general notion of independence, aptly named conditional independence.

Definition 3.4 (Conditional independence). Let (X, Σ_X) , (Y, Σ_Y) , (Z, Σ_Z) be measurable spaces, and let P be a joint probability measure on the product space $X \times Y \times Z$. We say that X and Y are *conditionally independent given Z* (with respect to P) if

- (i) a conditional $P_{XY|Z} : (Z, \Sigma_Z) \rightarrow (X \times Y, \Sigma_X \otimes \Sigma_Y)$ exists; and
- (ii) for each $z \in Z$, X and Y are independent with respect to the probability measure $P_{XY|Z}(z, -)$.

This notion gives us far more resolution in investigations of how variables can depend on each other, and hence in finding causal relationships. For example, the variables representing the amount of rain on a given day in London and in Beijing are dependent—on a winter day it on average rains more than a summer one in both cities—, but we can tell they are not causally related because they are conditionally independent given the season. A key feature of Bayesian networks is that it allows us to translate facts about causal relationships into facts about conditional independence, and vice versa.

The following lemma helps with further conceptualising conditional independence. In particular, conditions (iii) and (iv) say that if X and Y are conditionally independent given Z , then upon knowing the outcome of Z , the outcome of X gives no information about the outcome of Y , and the outcome of Y gives no information about the outcome of X .

Lemma 3.5 (Countable conditional independence). *Let X, Y, Z be countable discrete measurable spaces with a joint probability measure P such that the marginals on Z, XZ, YZ each have full support. The following are equivalent:*

- (i) X and Y are conditionally independent given Z .
- (ii) $P_{X|Z}(z, \{x\})P_{Y|Z}(z, \{y\}) = P_{XY|Z}(z, \{(x, y)\})$ for all $x \in X, y \in Y, z \in Z$.
- (iii) $P_{X|YZ}(y, z, \{x\}) = P_{X|Z}(z, \{x\})$ for all $x \in X, y \in Y, z \in Z$.
- (iv) $P_{Y|XZ}(x, z, \{y\}) = P_{Y|Z}(z, \{y\})$ for all $x \in X, y \in Y, z \in Z$.

Proof. The equivalence of (i) and (ii) is just the definition of conditional independence, noting that in the discrete case the conditionals are uniquely determined by their values on individual outcomes. The equivalence of (ii), (iii), and (iv) follow from elementary facts in probability theory; a proof can be found in [12]. \square

3.2 Bayesian networks

We introduce Bayesian networks with an example.


Example 3.6. Suppose that we wish to add a causal interpretation to a joint probability measure on the binary random variables $A = \{a, \neg a\}$, $B = \{b, \neg b\}$, and $C = \{c, \neg c\}$ representing the propositions that, upon being presented with a food:

- A: you like, or *appreciate*, the food.
- B: the food is nutritionally *beneficial*.
- C: you *choose* to eat the food.

Let these random variables have joint probability measure given by the table

x	$P(x)$
a, b, c	0.24
$a, b, \neg c$	0
$a, \neg b, c$	0.18
$a, \neg b, \neg c$	0.18
$\neg a, b, c$	0.06
$\neg a, b, \neg c$	0.10
$\neg a, \neg b, c$	0
$\neg a, \neg b, \neg c$	0.24

Intuitively, the causal relationships between our variables are obvious: liking a food influences whether you choose to eat it, and so does understanding it has health benefits, but otherwise there are no causal relationships between the variables—liking a food does not cause it to be more (or less) healthy. We shall represent these causal relationships by the directed graph



where we have drawn an arrow from one variable to another to indicate that that variable has causal influence on the other. The above joint probability measure and graph comprise what we will later define as a Bayesian network.

Note that we could not have chosen just any directed graph with vertices A , B , and C , as assertions about causal relationships have consequences that must be reflected in the joint probability measure. For example, as in the above graph neither A or B cause of each other, nor have a common cause, we expect that A and B are independent with respect to the marginal P_{AB} . This is true. Writing probability measures as $n \times 1$ stochastic matrices with respect to the bases $\{x, \neg x\}$ for the binary variables and $\{ab, a\neg b, \neg ab, \neg a\neg b\}$ for the variable AB , we have

$$P_A = \begin{pmatrix} 0.6 \\ 0.4 \end{pmatrix}; \quad P_B = \begin{pmatrix} 0.4 \\ 0.6 \end{pmatrix}$$

and hence

$$P_A \otimes P_B = P_{AB} = \begin{pmatrix} 0.24 \\ 0.36 \\ 0.16 \\ 0.24 \end{pmatrix}$$

Furthermore, the above graph suggests that the probability measure on C can be written as a function of the outcomes of both its causes A and B . We thus expect

that the measure has factorisation

$$P(x_a, x_b, x_c) = \int_{(x_a, x_b)} \int_{x_c} dP_{C|AB} dP_{AB} = \int_{x_a} \int_{x_b} \int_{x_c} dP_{C|AB} dP_B dP_A,$$

where $x_a \in A$, $x_b \in B$, $x_c \in C$. As these variables are finite, we might also write this requirement as

$$P(x_a, x_b, x_c) = P_{C|AB}(x_a, x_b, \{x_c\})P_B(\{x_b\})P_A(\{x_a\}).$$

Again, this is also true, with

$$P_{C|AB} = \begin{pmatrix} 1 & 0.5 & 0.375 & 0 \\ 0 & 0.5 & 0.625 & 1 \end{pmatrix}.$$

Motivated by this, we will later define a compatibility requirement in terms of the existence of a certain factorisation. Note that although in general a probability measure will have many factorisations; the directed graph specifies a factorisation that we attach greater—causal—significance to.

An advantage of expressing the causal relationships as a graph is that we may read from it other, acausal, dependencies. For example, while A and B are independent, the above graph suggests that if we know something about their common consequence C , this should induce some dependence between them. Indeed we find this does occur. Observe that

$$P_{AB|C}(c, -) = \begin{pmatrix} 0.5 \\ 0.375 \\ 0.125 \\ 0 \end{pmatrix},$$

which indicates that of the foods you choose to eat, the foods you like are more likely to be unhealthy than those you dislike.

More than this, however, marking certain relationships as causal affects our understanding of how a joint probability measure should be interpreted; we will see an example of this in the next chapter.

To make these ideas precise we introduce some definitions. Recall that a *directed graph* $G = (V, A, s, t)$ consists of a finite set V of vertices, a finite set A of arrows, and source and target maps $s, t : A \rightarrow V$ such that no two arrows have the same source and target—precisely, such that for all $a, a' \in A$ either $s(a) \neq s(a')$ or $t(a) \neq t(a')$. An arrow $a \in A$ is said to be an *arrow from u to v* if $s(a) = u$ and $t(a) = v$, while a sequence of vertices $v_1, \dots, v_k \in V$ is said to form a *path from v_1 to v_k* if for all $i = 1, \dots, k - 1$ there exists an arrow $a_i \in A$ from v_i to v_{i+1} . A path is also called a *cycle* if in addition $v_1 = v_k$. A directed graph is *acyclic* if it contains no cycles.

As demonstrated in the above example, directed acyclic graphs provide a depiction of causal relationships between variables; the direction represents the asymmetry of the causal relationship, while cycles are disallowed as variables cannot have causal influence on themselves. When we think of the set of vertices of a directed acyclic graph as the set of random variables of a system, we will also call the graph a *causal structure*.

Given a directed graph, we use the terminology of kinship to talk of the relationships between vertices, saying that a vertex u is a *parent* of a vertex v if there is an arrow from u to v , u is an *ancestor* of v if there is a path from u to v , u is a *child* of v if there is an arrow from v to u , and u is a *descendent* of v if there is a path from u to v . We will in particular talk of the parents of a vertex frequently, and so introduce the notation

$$\text{pa}(v) = \{u \in V \mid \text{there exists } a \in A \text{ such that } s(a) = u, t(a) = v\}$$

for the set of parents of a vertex. When dealing with graphs as causal structures, we will also use the names *direct causes*, *causes*, *direct effects*, and *effects* to mean parents, ancestors, children, and descendants respectively.

We say that an ordering $\{v_1, \dots, v_n\}$ of the set V is an *ancestral ordering* if v_i is an ancestor of v_j only when $i < j$.

Definition 3.7 (Bayesian network). Let $G = (V, A, s, t)$ be a directed acyclic graph, for each $v \in V$ let X_v be a measurable space, and let P be a joint probability measure on $\prod_{v \in V} X_v$. We say that the causal structure G and the joint probability measure P are *compatible* if there exists an ancestral ordering of the elements V such that there exist conditionals such that

$$P(A_1 \times A_2 \times \dots \times A_n) = \int_{A_1} \int_{A_2} \dots \int_{A_n} P_{X_n|\text{pa}(X_n)} \dots P_{X_2|\text{pa}(X_2)} P_{X_1}.$$

A *Bayesian network* (G, P) is a pair consisting of a compatible joint probability measure and causal structure.

A better understanding of this compatibility requirement can be gained from the examining following theorem.

Theorem 3.8 (Equivalent compatibility conditions). *Let G be a causal structure, let $\{(X_v, \Sigma_v) \mid v \in V\}$ be a collection of finite measurable spaces indexed by the vertices of G , and let P be a joint probability measure on their product space. Then the following are equivalent:*

- (i) P is compatible with G .
- (ii) $P(x_1, \dots, x_n) = \prod_{i=1}^n P_{X_i|\text{pa}(X_i)}(\text{pa}(x_j), \{x_i\})$, where $\text{pa}(x_j)$ is the tuple consisting of x_j such that $X_j \in \text{pa}(X_i)$.
- (iii) P obeys the ordered Markov condition with respect to G : given any ancestral ordering of the variables, each variable is independent of its remaining preceding variables conditional on its parents.
- (iv) P obeys the arental Markov condition with respect to G : each variable is independent of its nondescendents conditional on its parents.

Proof. This follows from Corollaries 3 and 4 to Theorem 3.9 in Pearl [18]. □

As an illustration of the relevance of causal structure, we note that conditional independence relations between variables of a Bayesian network can be read from the causal structure using a straightforward criterion. Call a sequence of vertices $v_1, \dots, v_k \in V$ an *undirected path* from v_1 to v_k if for all $i = 1, \dots, k - 1$ there exists an arrow $a_i \in A$ from v_i to v_{i+1} or from v_{i+1} to v_i . An undirected paths v_1, v_2, v_3 of three vertices then take the form of a

- (i) *chain*: $v_1 \rightarrow v_2 \rightarrow v_3$ or $v_1 \leftarrow v_2 \leftarrow v_3$,
- (ii) *fork*: $v_1 \leftarrow v_2 \rightarrow v_3$; or
- (iii) *collider*: $v_1 \rightarrow v_2 \leftarrow v_3$.

An undirected path from a vertex u to a vertex v is said to be *d-separated* by a set of nodes S if either the path contains a chain or a fork such that the centre vertex is in S , or if the path contains a collider such that the neither the centre vertex nor any of its descendants are in S . A set S then *d-separates* a set U from a set T if every path from a vertex in U to a vertex in T is *d-separated* by S .

The main idea of this definition is that causal influence possibly creates dependence between two random variables if one variable is a cause of the other, the variables have a common cause, or the variables have a common consequence and the outcome of this consequence is known. In this last case, knowledge of the consequence ‘unseparates’ the two variables along the path through the known common consequence. On the other hand, any information gained through having a common cause is rendered moot if we have knowledge about a variable through which the causal influence is mediated. These ideas are captured by the following theorem.

Theorem 3.9. *Let (G, P) be a Bayesian network. If sets U and T of vertices of G are d-separated by a third set S , then with respect to P , the product random variables $\prod_{u \in U} X_u$ and $\prod_{t \in T} X_t$ are conditionally independent given $\prod_{s \in S} X_s$.*

Proof. See Verma and Pearl [24]. □

An example of this is the way the variables A and B are independent in Example 3.6, but dependent conditional on C .

Chapter 4

Causal Theories

We now tie the elements of the last three chapters together to propose and develop a novel algebraic structure: a causal theory. After introducing these structures, we discuss their models in various categories and how such models might be interpreted, and then look at a possibly confusing situation that causal theories and their associated graphical language help make lucid.

4.1 The category associated to a causal structure

We wish to fashion a category that captures methods of reasoning with causal relationships. In this category, we will want our objects to represent the variables of a situation, while the morphisms should represent the ways one can deduce knowledge about one variable from another. Furthermore, as we will want to deal with more than one variable at a time, and the outcomes their joint variable may take, this category will be monoidal.

As we may only reason about causal relationships once we have some causal relationships to reason with, we start by fixing a set of symbols for our variables and the causal relationships between them. Let $G = (V, A, s, t)$ be a directed acyclic graph. From this we construct a strict symmetric monoidal category \mathcal{C}_G in the following way.


For the objects of \mathcal{C}_G we take the set \mathbb{N}^V of functions from V to the natural numbers. These may be considered collections of elements of the set of variables V , allowing multiplicities, and we shall often just write these as strings w of elements of V . Here the order of the symbols in the string is irrelevant, and we write \emptyset for empty string, which corresponds to the zero map of \mathbb{N}^V . We view these objects as the *variables* of the causal theory, and we further call the objects which are collections consisting of just one instance of a single element of V the *atomic variables* of the causal theory.

There are two distinct classes of generating morphisms for \mathcal{C}_G . The first class is the collection of *comonoid maps*: for each atomic variable $v \in V$, we include morphisms $\mathcal{Y}_v : v \rightarrow vv$ and $\mathcal{I}_v : v \rightarrow \emptyset$. These represent the ideas of duplicating some information about v , or forgetting some. The maps of the second class are called the *causal mechanisms*. These consist of, for each atomic variable $v \in V$, a morphism $[v|pa(v)] : pa(v) \rightarrow v$, where $pa(v)$ is the string consisting of the parents of v in any order, and represent the ways we may use information about a collection of variables to infer facts about another. We then use these morphisms as generators for a strict symmetric monoidal category, taking all products and well-defined compositions, subject only to the constraint that for each $v \in V$ the pair $(\mathcal{Y}_v, \mathcal{I}_v)$ forms a comonoid. As the swaps are identity maps, these comonoids are immediately commutative.


We call this category \mathcal{C}_G the *causal theory* of the causal structure G . Morphisms of \mathcal{C}_G represent ways to reason about the outcome of the codomain variable given some knowledge about the domain variable.

As the causal mechanisms are labelled with their domain and codomain, there is usually no need to label the strings when representing morphisms of \mathcal{C} with string diagrams. We also often do not differentiate between the comonoid maps with labels, as the context makes which comonoid map we are applying. The order in which we write the string representing the set $pa(v)$ corresponds to the order of the input strings.

Example 4.1. The causal theory of the causal structure



of Example 3.6 is the symmetric monoidal category with objects collections of the letters A , B , and C , and morphisms generated by counit and comultiplication maps on each of A , B , and C , as well as causal mechanisms $[A] : \emptyset \rightarrow A$, $[B] : \emptyset \rightarrow B$, and $[C|AB] : AB \rightarrow C$. We depict these causal mechanisms respectively as



We now list a few facts to give a basic understanding of the morphisms in these categories. These morphisms represent predictions of the consequences of the domain

variable on the codomain variable. As causal structures are acyclic—giving rise to a ‘causal direction’, or (noncanonical) ordering on the variables—, causal theories similarly have such a direction, and this puts limits on the structure. Indeed, one consequence is that a morphism can only go from an effect to a cause if it factors through the monoidal unit; this represents ‘forgetting’ the outcome of the effect, and reasoning about the outcomes of the cause from other, background, information. We say a map is *inferential* if it does not factor through the monoidal unit.

Proposition 4.2. *Let \mathcal{C}_G be a causal theory, and let v, v' be atomic variables in \mathcal{C}_G . If there exists an inferential map $v \rightarrow v'$, then v is an ancestor of v' in G .*

Proof. We reason via string diagrams to prove the contrapositive.

Observe that a generating map is inferential if and only if, in its string diagram representation, the domain is topologically connected to the codomain, and that this property is preserved by the counitality relation the comonoid maps must obey. Thus it is also true in general: a morphism in \mathcal{C}_G is inferential if and only if, in all string diagram representations, the domain is topologically connected to the codomain.

Note also that for all generating maps with string diagrams in which the domain and codomain are connected, the domain and codomain are nonempty and each element of the domain is either equal to or an ancestor of each element of the codomain. This property is also preserved by the counitality relation. Thus, if v is not an ancestor of v' , in all string diagram representations of a map $v \rightarrow v'$ the domain is not topologically connected to the codomain. Taking any such string diagram and continuously deforming it by moving all parts of the component connected to the domain below all parts of the component connected to the codomain, we thus see that the map may be rewritten as one that factors through the monoidal unit. \square

In fact the converse also holds: if v is an ancestor of v' , then there always exists an inferential map $v \rightarrow v'$. Indeed, if w, w' are objects in \mathcal{C}_G containing each atomic variable no more than once—that is, when $w, w' \subseteq V$ —, we can construct a map $w \rightarrow w'$ in the following manner.

1. Take the smallest subgraph $G_{w \rightarrow w'}$ of G containing the vertices of w and w' , and all paths in V that terminate at an element of w' and do not pass through w .¹

¹Precisely, this means we take the subgraph with set of arrows $A_{w \rightarrow w'} \subseteq A$ consisting of all $a \in A$ for which there exist a_1, \dots, a_n with

(i) $s(a_1) = t(a)$,

2. For each vertex $v \in V_{w \rightarrow w'}$ let k_v be the number of arrows of $A_{w \rightarrow w'}$ with source v . Then:
 - (i) for each $v \in w$ take the string diagram representing the composition of $k_v - 1$ comultiplications on v or, when $k_v = 0$, the counit of v .
 - (ii) For each $v \in w'$ take the string diagram for $[v|pa(v)]$ composed with a sequence of k_v comultiplications on v .
 - (iii) For each $v \in V_{w \rightarrow w'} \setminus (w \cup w')$ take the string diagram for $[v|pa(v)]$, composed with either a sequence of $k_v - 1$ comultiplications on v or, if $k_v = 0$, composed with the counit on v .
3. From this collection of string diagrams, create a single diagram by connecting an output of one string diagram to an input of another if the set $A_{w \rightarrow w'}$ contains an arrow from the indexing vertex of the first diagram to the indexing vertex of the second.

Due to the symmetry of the monoidal category \mathcal{C}_G and the associativity of the comultiplication maps, this process uniquely defines a string diagram representing a morphism from w to w' . Moreover, this map is inferential whenever there exists some $v \in w$ that is an ancestor of some $v' \in w'$. These maps are in a certain sense the uniquely most efficient ways of predicting probabilities on w' using information about w , and will play a special role in what follows. For short, we will call these maps *causal conditionals* and write them maps $[w'|w]$, or simply $[w']$ when $w = \emptyset$. In this last case, we will also call the map $[w']$ the *prior on* $[w']$.


Example 4.3. This construction is a little abstruse on reading, but the main idea is simple and an example should make it much clearer. Let G be the causal structure

-
- (ii) $s(a_{i+1}) = t(a_i)$ for $i = 1, \dots, n - 1$,
 - (iii) $t(a_n) \in w'$,
 - (iv) $s(a_i) \notin w$ for $i = 1, \dots, n$,


and vertices

$$V_{w \rightarrow w'} = w \cup w' \cup \{v \in V \mid v = s(a) \text{ or } t(a) \text{ for some } a \in A_{w \rightarrow w'}\}.$$

Note that for each $v \in V_{w \rightarrow w'} \setminus w$, the set of parents of v in this subgraph is equal to the set of parents of v in the whole graph G .




and suppose that we wish to compute the causal conditional $[DE||B]$. Step 1 gives the subgraph $G_{B \rightarrow DE}$




consisting of all paths to D or E not passing through B .

Step 2 then states that the causal conditional $[DE||B]$ comprises the maps



and we then compose these mimicking the topology of the graph $G_{B \rightarrow DE}$ to give the map



4.2 Interpretations of causal theories

Causal theories express abstractly avenues of causal reasoning, but this serves no purpose in describing specific causal relationships until we attach meanings, or an interpretation, to the objects and morphisms of the theory. The strength of separating out the syntax of reasoning is that these interpretations may now come from any symmetric monoidal category. Stated formally, let \mathcal{C} be a causal theory, and let \mathcal{D} be any symmetric monoidal category. Then a *model of \mathcal{C} in \mathcal{D}* , or just a *causal model*, is a strong monoidal functor $M : \mathcal{C} \rightarrow \mathcal{D}$.

We explore the basic properties of causal models in a few categories. To demonstrate the basic ideas, we first take a brief look at models in **Set** and **Rel**; models in these categories will be useful for describing deterministic and possibilistic causal relationships respectively. While **Meas** is another obvious candidate setting for examining causal models, we merely note that causal models here behave somewhat similarly to **Set** and move on to **Stoch**, the main category of interest. Here causal models generalise Bayesian networks. As Bayesian networks are known to provide a useful tool for the discussion of causality, this lends support to the idea that the richer structure of causal models in **Stoch** do too.

Models in Set

Due to its familiarity, we begin our discussion of causal models with an examination of the forms they take in **Set**. In both **Set** and **Rel** the objects are sets. For the purposes of causal models, it is useful to view these sets as variables, with the elements the possible outcomes of the variable. With this interpretation, we can understand **Set** as a subcategory of **Meas** in which every measurable space is discrete, making it possible to measure any subset of the outcomes of each variable. Morphisms in **Set**—that is, set functions—then assign a single outcome of the codomain variable to each outcome of the domain variable, and so can be said to describe deterministic causal relationships.

Given a causal theory \mathcal{C} , a model of \mathcal{C} in **Set** by definition consists of a strong monoidal functor $M : \mathcal{C} \rightarrow \mathbf{Set}$. To specify such a functor up to isomorphism, it is enough to specify the image of each atomic variable and each generating map, subject to the constraints that the generating maps chosen are well-typed with respect to the chosen images of the atomic variables, and that the images of the comultiplication and counit obey the laws of a commutative comonoid. Indeed, once these are specified, the values of the functor on the remaining objects and morphisms of \mathcal{C} are, up to

isomorphism, determined by the definition of a strong monoidal functor. Note also that as long as the aforementioned constraints are fulfilled we have a well-defined strong monoidal functor.

We first observe, as we will also in the case of **Stoch**, that each object of **Set** has a unique comonoid structure, and this comonoid is commutative. To wit, for each set X , there is a unique map $X \rightarrow \{*\}$, taking the product of this map and the identity map $X \rightarrow X$ gives the projection maps $X \times X \rightarrow X$, and the only function $X \rightarrow X \times X$ that composes to the identity with the projection map on each factor is the diagonal map $x \mapsto (x, x)$. Moreover, choosing the diagonal map as a comultiplication indeed gives a commutative comonoid with this map. It is a consequence of this that we need not worry about the comonoid maps; choosing a set for each variable also chooses the comonoid maps for us.

On the other hand, as the causal mechanisms need not obey any equations, so having defined a map on the objects, any choice of functions from the product set of all the direct causes of each variable to the variable itself then gives a model of the causal theory in **Set**. Each such function returns the outcome of its codomain variable given a configuration of the outcomes of its causes. In this sense a model of a causal theory in **Set** specifies how causes affect their consequences in a deterministic way.

As maps from the monoidal unit in **Set** are just a pointings of the target set, the priors $M[w]$ of a causal model are just a choice of an outcome for each of the atomic variables in w . In the case of an atomic variable with no causes, the prior $M[v] : \{*\} \rightarrow Mv$ is simply the causal mechanism $M[v|pa(v)]$, and just picks an element of the set Mv . One might interpret this as the ‘default’ state of the variable v , and subsequently interpret the prior $M[V]$ on the set of all variables V as the default state of all variables in the system.

We shall see this as a general feature of models of causal theories; the priors specify what can be known about the system in some default state, while more generally the morphisms describe the causal relationships between variables even when not in this state.


Models in Rel

In the category **Rel**, we interpret a relation $r : X \rightarrow Y$ to mean that that if X has the outcome $x \in X$, then Y may only take the outcomes y related to x via r . This is a possibilistic notion of causality, in which the outcomes of the causes do not determine

a single outcome of the effect variable as in **Set**, but only put some constraint on the possible outcomes of the effect variable.

A curious property of the category **Rel** is that any relation $X \rightarrow Y$ may also be viewed as a relation $Y \rightarrow X$ in a natural way—that is, **Rel** is equipped with a contravariant endofunctor that squares to the identity, or a dagger functor. This means that we have a way of reversing the direction any morphism we choose, and in this sense **Rel** itself is acausal. This makes causal models all the more useful when working in **Rel**, as they provide a way of privileging certain relations with a causal direction.

For any object in **Rel**, we may view the functions forming the unique comonoid on this set in **Set** as relations, and hence have a commutative comonoid in **Rel**. An interesting collection of causal models $M : \mathcal{C} \rightarrow \mathbf{Rel}$ in **Rel** are those in which all objects are given these comonoid structures. Note that a map $s : \{*\} \rightarrow X$ from the monoidal unit to an object X in **Rel** is simply a subset S of X and, assuming X has the comonoid structure in which the comonoid maps are functions, for any relation $r : X \rightarrow Y$ the composite



is then equal to the set

$$\{(x, y) \in X \times Y \mid x \in S, y \sim_r x\}.$$

Thus when the comonoid maps of the model are those of **Rel** inherited from **Set**, it is easy to see that the priors $[w]$ in **Rel** are given by the subset of the product set of the atomic variables in w consisting of all joint outcomes that are possible given the constraints of the causal mechanisms.

In this setting, however, commutative comonoids are more general; for example, any collection of abelian groups forms a commutative comonoid on the union of the sets of elements of these groups [16]. We leave examination of causal structures for these other comonoid structures, and their interpretations, for later work.

Models in Stoch

We begin our discussion of causal models in **Stoch** by showing that for **Stoch** too we need not worry about selecting comonoid maps; the deterministic comonoid structure **Stoch** inherits from **Set** is the only comonoid structure on each object.

Lemma 4.4. *Each object of **Stoch** has a unique comonoid structure. Moreover, this comonoid structure is commutative.*

Proof. Fix an object (X, Σ) in **Stoch**. We first show the existence of a comonoid structure on X by showing that the stochastic maps

$$\Psi : (X, \Sigma) \longrightarrow (X \times X, \Sigma \otimes \Sigma);$$

defined by

$$\begin{aligned} \Psi : X \times (\Sigma \otimes \Sigma) &\longrightarrow [0, 1]; \\ (x, A) &\longmapsto \begin{cases} 1 & \text{if } (x, x) \in A, \\ 0 & \text{if } (x, x) \notin A, \end{cases} \end{aligned}$$

and

$$\dagger : (X, \Sigma) \longrightarrow (*, \{\emptyset, *\})$$

defined by

$$\begin{aligned} \dagger : X \times \{\emptyset, *\} &\longrightarrow [0, 1]; \\ (x, *) &\longmapsto 1, \\ (x, \emptyset) &\longmapsto 0, \end{aligned}$$

form the comultiplication and counit for a comonoid structure respectively.

Indeed, observe that both these maps are deterministic stochastic maps, with Ψ specified by the measurable function $X \rightarrow X \times X; x \mapsto (x, x)$ and \dagger specified by $X \rightarrow *; x \mapsto *$. From here it is straightforward to verify that these functions obey coassociativity and counitality as functions in **Meas**, and hence these identities are true in **Stoch**.

We next prove uniqueness. As the monoidal unit $(*, \{\emptyset, *\})$ is terminal, there is a unique stochastic map from (X, Σ) to the terminal object, and so \dagger is the only possible choice for the counit of a comonoid on (X, Σ) . Suppose that $\delta : (X, \Sigma) \rightarrow (X \times X, \Sigma \otimes \Sigma)$ is a stochastic map such that (δ, \dagger) forms a comonoid. We will show in fact that for all $x \in X$ and $A \in \Sigma \otimes \Sigma$ we have

$$\delta(x, A) = \begin{cases} 1 & \text{if } (x, x) \in A; \\ 0 & \text{if } (x, x) \notin A, \end{cases}$$

and so $\delta = \Psi$. Note that as $\delta_x = \delta(x, -)$ is a probability measure on $(X \times X, \Sigma \otimes \Sigma)$ it is enough to show that $\delta(x, A) = 1$ whenever $(x, x) \in A$.

To begin, note that counitality on the right implies that

$$\begin{array}{c} \text{---} \\ | \\ \boxed{\delta} \\ | \\ \text{---} \end{array} = \begin{array}{c} \bullet \\ | \\ \text{---} \end{array}$$

so for all $x \in X$ and $B \in \Sigma$ we have

$$\begin{aligned} \delta_x(B \times X) &= \delta(x, B \times X) = \int_{X \times X} \chi_{B \times X} d\delta_x \\ &= \int_{X \times X} \mathbb{1}(-, -, B) d\delta_x \\ &= \begin{array}{c} \bullet \\ | \\ \boxed{\delta} \\ | \\ \text{---} \end{array}(x, B) \\ &= \text{id}_{(X, \Sigma)}(x, B) \\ &= \begin{cases} 1 & \text{if } x \in B; \\ 0 & \text{if } x \notin B. \end{cases} \end{aligned}$$

Similarly, counitality on the left implies that for all $x \in X$ and $B \in \Sigma$ we have

$$\delta_x(X \times B) = \begin{cases} 1 & \text{if } x \in B; \\ 0 & \text{if } x \notin B. \end{cases}$$

We shall use these facts in the following.

Fix $x \in X$ and let now $A \in \Sigma \otimes \Sigma$ be such that $(x, x) \in A$. Recalling our characterisation of product σ -algebras in Example 2.4, we may assume A is of the form $A = \cup_{i \in I} (C_i \times D_i)$, where I is a countable set and $C_i, D_i \in \Sigma$ for all $i \in I$. There thus exists $0 \in I$ such that $(x, x) \in C_0 \times D_0$, and hence $x \in C_0$ and $x \in D_0$. Since we have shown above that this implies that $\delta_x(C_0^c \times X) = \delta_x(X \times D_0^c) = 0$, we then have

$$\begin{aligned} \delta_x(A) &= \delta_x(\cup_{i \in I} (C_i \times D_i)) \\ &\geq \delta_x(C_0 \times D_0) \\ &= \delta_x((C_0 \times X) \cap (X \times D_0)) \\ &= 1 - \delta_x((C_0^c \times X) \cup (X \times D_0^c)) \\ &\geq 1 - (\delta_x(C_0^c \times X) + \delta_x(X \times D_0^c)) \\ &= 1, \end{aligned}$$

so $\delta(x, A) = 1$ as required.

It remains to check that this comonoid is commutative. Recalling that the swap $\times : (X \times X, \Sigma \otimes \Sigma) \rightarrow (X \times X, \Sigma \otimes \Sigma)$ on $(X, \Sigma) \otimes (X, \Sigma)$ is the deterministic stochastic map given by $X \times X \rightarrow X \times X; (x, y) \mapsto (y, x)$, it is immediately clear that the comultiplication \vee is commutative. \square

Arguing as for models of **Set**, given a causal theory \mathcal{C} , strong monoidal functors $P : \mathcal{C} \rightarrow \mathbf{Stoch}$ are thus specified by arbitrary choices of measurable space for each atomic variable, and a subsequent arbitrary choices of causal mechanisms of the required domain and codomain.

These interpretations of the causal mechanisms give rise to a joint probability measure compatible with the causal structure underlying the causal theory. Indeed, this can be seen as the key difference between a model of a causal theory in **Stoch** and a Bayesian network: models of causal theories privilege factorisations, while Bayesian networks only care about the joint probability measure.

Theorem 4.5. *Let G be a directed acyclic graph with vertex set V and let $P : \mathcal{C}_G \rightarrow \mathbf{Stoch}$ be a model of the causal theory \mathcal{C}_G in **Stoch**. Then the causal structure G and the probability measure defined by the prior $P[V]$ are compatible.*

Proof. Recall that $[V]$ is the prior on the collection consisting of one copy of each of the atomic variables V . For each $v \in V$ we have a measurable space Pv , and as $P[V] : \emptyset \rightarrow PV$ is a point of **Stoch** it defines a joint probability measure on the product measurable space PV . We must show that $P[V]$ has the required factorisation.

To this end, choose some ancestral ordering of V , writing V now as $\{v_1, \dots, v_n\}$ with the elements numbered according to this ordering. By construction, the string diagram of the prior $[V]$ consists of one copy of each causal mechanism $[v_i | \text{pa}(v_i)]$ and k_i copies of each comultiplication \cup_{v_i} , where k_i is the number of children of the vertex v_i . As each v_i appears exactly once as the codomain of the causal mechanisms $[v_i | \text{pa}(v_i)]$, the coassociativity of each comonoid and the rules of the graphical calculus for symmetric monoidal categories, show that any way of connecting these elements to form a morphism $\emptyset \rightarrow V$ produces the same morphism. In particular, we may build $[V]$ as the composite of the morphisms $[V]_i : v_1 \dots v_{i-1} \rightarrow v_1 \dots v_i$ defined by

$$[V]_i = \begin{array}{c} \begin{array}{c} \dots \\ | \\ \dots \end{array} \quad \begin{array}{c} \dots \\ | \\ \dots \end{array} \quad \begin{array}{c} v_1 v_2 \dots v_i \\ \cup \quad \cup \\ \dots \quad \dots \\ | \quad | \\ \dots \quad \dots \\ v_1 v_2 \dots v_{i-1} \end{array} \\ (v_j \notin \text{pa}(v_i)) \quad (v_j \in \text{pa}(v_i)) \end{array}$$

In words, the morphism $[V]_i$ is the morphism $v_1 \dots v_{i-1} \rightarrow v_1 \dots v_i$ constructed by applying a comultiplication to each of the parents of v_i , and then applying the causal morphism $[v_i|\text{pa}(v_i)]$. Note that as we have ordered the set V with an ancestral ordering, all parents of v_i do lie in the set of predecessors of v_i .

Observe now that given any stochastic map $k : (X, \Sigma_X) \rightarrow (Y, \Sigma_Y)$, if Υ is the unique comultiplication on X , then the composite

$$\begin{array}{c} \text{---} \\ | \\ \text{---} \\ | \\ \text{---} \end{array} \begin{array}{c} \boxed{k} \\ | \\ \text{---} \end{array} : (X, \Sigma_X) \longrightarrow (X \times Y, \Sigma_X \otimes \Sigma_Y)$$

is given by

$$((\text{id}_X \otimes k) \circ \Upsilon)(x, A \times B) = \int_{X \times X} \chi_A(-) k(-, B) d\Upsilon_x = \chi_A(x) k(x, B)$$

for all $x \in X$, $A \in \Sigma_X$, $B \in \Sigma_Y$. Furthermore, if $\mu : * \rightarrow (X, \Sigma_X)$ is a measure on X , then

$$\begin{array}{c} \text{---} \\ | \\ \text{---} \\ | \\ \text{---} \end{array} \begin{array}{c} \boxed{k} \\ | \\ \text{---} \end{array} : (*, \{\emptyset, *\}) \longrightarrow (X \times Y, \Sigma_X \otimes \Sigma_Y)$$

$$\begin{array}{c} \text{---} \\ | \\ \text{---} \\ | \\ \text{---} \end{array} \begin{array}{c} \mu \\ | \\ \text{---} \end{array}$$

is given by

$$((\text{id}_X \otimes k) \circ \Upsilon \circ \mu)(A \times B) = \int_A k(x, B) d\mu$$

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

Thus, taking the image under P , each of the maps $[V]_i$ gives

$$P[V]_i(x_1, \dots, x_{i-1}, A_1 \times \dots \times A_i) = \prod_{j=1}^{i-1} \chi_{A_j}(x_j) P[v_i|\text{pa}(v_i)](x_1, \dots, x_{i-1}, A_i).$$

for all $x_j \in P v_j$, $A_j \in \Sigma_{P v_j}$, and composing them gives

$$\begin{aligned} P[V](A_1 \times \dots \times A_n) &= P[V]_n \circ \dots \circ P[V]_2 \circ P[V]_1(A_1 \times A_2 \times \dots \times A_n) \\ &= \int_{A_1} P[V]_n \circ \dots \circ P[V]_2(-, A_2 \times \dots \times A_n) dP[v_1|\text{pa}(v_1)] \\ &= \quad \quad \quad \vdots \\ &= \int_{A_1} \int_{A_2} \dots \int_{A_n} dP[v_n|\text{pa}(v_n)] \dots dP[v_2|\text{pa}(v_2)] dP[v_1|\text{pa}(v_1)] \end{aligned}$$

for all $A_j \in \Sigma_{P v_j}$. This is a factorisation of $P[V]$ of the required type. \square

We call the pair $(G, P[V])$ the *Bayesian network induced by P* . Thus we see that, given a causal structure and any stochastic causal model of its theory, the induced joint distribution on the atomic variables of the theory forms a Bayesian network with the causal structure. On the other hand, if we have a Bayesian network on this causal structure, we may construct a stochastic causal model inducing this distribution, but only by picking some factorisation of our joint distribution. To iterate, this is the key distinction between Bayesian networks and stochastic causal models: a Bayesian network on a causal structure requires only that there *exist* a factorisation for the distribution respecting the causal structure, while a stochastic causal model *explicitly chooses* a factorisation.

An advantage of working within a causal theory, rather than just with the induced Bayesian network, is that the additional structure allows neat representations of operations that one might want to do to a Bayesian network. The remainder of this dissertation comprises a brief exploration of this. We conclude this section by noting that the priors of the causal theory represent the marginals of the induced Bayesian network.

Theorem 4.6. *Given a model of $P : \mathcal{C} \rightarrow \mathbf{Stoch}$ of a causal theory \mathcal{C}_G and a set w of atomic variables of \mathcal{C}_G , the prior $P[w]$ is equal to the marginal on the product measurable space of the variables in w of the induced Bayesian network.*

Proof. We first note a more general fact: given a joint probability measure on $X \times Y$ expressed as a point in \mathbf{Stoch} , marginalisation over Y can be expressed as the composite of this point with the product of the identity \mid on X and counit \uparrow on Y . Indeed, if $\mu : * \rightarrow X \times Y$ is a probability measure, then

$$(\uparrow \circ \mu)(A) = \int_{X \times Y} \uparrow(x, y, A) d\mu = \int_{X \times Y} \chi_A(x) d\mu = \mu_X(A).$$

Thus the marginals of $P[V]$ may be expressed by composing $P[V]$ with counits on the factors marginalised over. We wish to show that these are the priors $P[w]$ of \mathcal{C} . Reasoning inductively, to show this it is enough to show that the composite of a prior with the product of a counit on one of its factors and identity maps on the remaining factors is again a prior.

Let w be a set of atomic variables of \mathcal{C} and let $v \in w$. We will show that the composite of $P[w]$ with the product of the counit on v and identity on $w \setminus \{v\}$ is equal to the prior $P[w \setminus \{v\}]$. We split into two cases: when v has a consequence in $G_{\emptyset \rightarrow w}$, and when v has no consequences in $G_{\emptyset \rightarrow w}$. For the first case, observe that $G_{\emptyset \rightarrow w} = G_{\emptyset \rightarrow w \setminus \{v\}}$ and $k_v \geq 1$. Thus the priors $[w]$ and $[w \setminus \{v\}]$ are the same but

for the fact we compose with one extra comultiplication after the causal mechanism $[v|\text{pa}(v)]$ in the case of the prior $[w]$. Thus the composite of $[w]$ with a counit on v is equal to $[w \setminus \{v\}]$ by the counitality law on v .

To deal with the second case we must work in **Stoch** and make use of the fact that the monoidal unit in **Stoch** is terminal. Indeed, as the monoidal unit in **Stoch** is terminal, in **Stoch** we have the equality of morphisms

$$\begin{array}{c} \bullet \\ | \\ \boxed{v|\text{pa}(v)} \\ | \dots | \\ \text{pa}(v) \end{array} = \begin{array}{c} \bullet \quad \bullet \\ | \quad | \\ \dots \\ \text{pa}(v) \end{array}$$

As v has no consequences in $G_{\emptyset \rightarrow w}$, the causal mechanism $[v|\text{pa}(v)]$ is not followed with any comultiplications in the construction of $[w]$. Thus, after composing $P[w]$ with a counit on v we may invoke the above identity, and then invoke the counitality law for each of the parents of v . This means that $P[w]$ is equal to the morphism constructed without the causal mechanism $[v|\text{pa}(v)]$, and with one fewer comultiplications on each of the parents of v . But this is precisely the morphism $P[w \setminus \{v\}]$. This proves the theorem. \square

4.3 Application: visualising Simpson’s paradox

One of the strengths of causal theories is that their graphical calculi provide a guide to which computations should be made if one wants to respect a causal structure, and in doing so also clarify what these computations mean. An illustration can be found in an exploration of confounding variables and Simpson’s paradox. This section owes much to Pearl [19, Chapter 6], extending the basics of that discussion with our new graphical notation.

Simpson’s paradox refers to the perhaps counterintuitive fact that it is possible for to have data such that, for all outcomes of a confounding variable, a fixed outcome of the independent variable makes another fixed outcome of the dependent variable *more* likely, and yet also that upon aggregation over the confounding variable the same fixed outcome of the independent variable makes the same fixed outcome of the dependent variable *less* likely. This is perhaps best understood through an example.

Consider the following, somewhat simplified, scenario: let us imagine that we wish to test the efficacy of a proposed new treatment for a certain heart condition. In our clinical experiment, we take two groups of patients each suffering from the heart condition and, after treating the patients in the proposed way, record whether

they recover. In addition, as we know that having a healthy blood pressure is also an important factor in recovery, we also take records of whether the blood pressure of the patient is within healthy bounds or otherwise at the conclusion of the treatment programme. This gives three binary variables: an independent variable $T = \{t, \neg t\}$, a dependent variable $R = \{r, \neg r\}$, and a third, possibly confounding variable $B = \{b, \neg b\}$, where we will think of the variables as representing the truth or otherwise of the following propositions:

T: the patient receives *treatment* for heart condition.

R: the patient has *recovered* at the conclusion of treatment.

B: the patient has healthy *blood pressure* at post-treatment checkup.

Suppose then that our experiment yields the data of Figure 4.1.


		T	
		t	$\neg t$
R	r	39	42
	$\neg r$	61	58

		TB			
		t, b	$\neg t, b$	$t, \neg b$	$\neg t, \neg b$
R	r	30	40	9	2
	$\neg r$	10	40	51	18

Figure 4.1: The table on the left displays the experiment data for the treatment and recovery of patients, while the table on the right displays the experiment when further subdivided according to the blood pressure of the patients measured post-treatment.


In these data we see the so-called paradox: for both patients with healthy and unhealthy blood pressure, treatment seems to significantly improve the chance of recovery, with the recovery rates increasing from 50% to 80% and from 10% to 15% respectively when patients are treated. On the other hand, when the studies are taken a whole, it seems treatment has no significant effect on the recovery rate, which drops slightly from 42% to 39%. Given this result, it is not clear whether the experiment indicates that treatment improves or even impairs chance of recovery. Should we or should we not then recommend the treatment?

The answer depends on the causal relationships between our variables. Suppose that the treatment acts in part via affecting blood pressure. Then the causal structure of the variables is given by the graph



In this case we should make our decision with respect to the aggregated data: else when we condition on the post-treatment blood pressures we eliminate information about how the treatment is affecting blood pressure, and so eliminate information about an important causal pathway between treatment and recovery. We therefore should not recommend treatment—although when we control for blood pressure the treatment seems to improve chances of recovery, the treatment also makes it less likely that a healthy blood pressure will be reached, offsetting any gain.

On the other hand, suppose that the treatment works in a way that has no effect on blood pressure. Then from the fact that the blood pressure and treatment variables are not independent we may deduce that the blood pressure variable biased selection for the treatment trial, and so the causal structure representing these variables is



Here we *should* pay attention to the data when divided according to blood pressure, as by doing this we control for the consequences of this variable. We then see that no matter whether a patient has factors leading to healthy or unhealthy blood pressure, the treatment raises their chance of recovery by a significant proportion.

These ideas are codified in the corresponding causal theories and their maps, with the causal effect of treatment on recovery expressed via the causal conditional $[R||T]$. For the first structure, let the corresponding causal theory be \mathcal{C}_1 , and the data give the following interpretations of the causal mechanisms $P : \mathcal{C}_1 \rightarrow \mathbf{FinStoch}$:

$$\begin{array}{c} \downarrow \\ T \end{array} = \begin{pmatrix} 0.5 \\ 0.5 \end{pmatrix} \quad \begin{array}{c} \boxed{B|T} \\ \downarrow \end{array} = \begin{pmatrix} 0.4 & 0.8 \\ 0.6 & 0.2 \end{pmatrix}$$

$$\boxed{R|TB} = \begin{pmatrix} 0.75 & 0.15 & 0.5 & 0.1 \\ 0.25 & 0.85 & 0.5 & 0.9 \end{pmatrix}$$

Here we have written the maps as their representations in **SMat** with respect to the basis ordering given in our definition of the variables. The causal conditional $P[R||T]$ is then

$$P[R||T] = \begin{array}{c} R \\ | \\ \boxed{R|TB} \\ | \\ \cup \\ \boxed{B|T} \\ | \\ T \end{array} = \begin{pmatrix} 0.39 & 0.42 \\ 0.61 & 0.58 \end{pmatrix}$$

The elements of the first row of this matrix represent the probability of recovery given treatment and no treatment respectively, and so this agrees with our assertion that in this case we should view the treatment as ineffective, and perhaps marginally harmful.

On the other hand, writing the corresponding causal theory to the second causal structure as \mathcal{C}_2 , in this case the data gives the stochastic model $Q : \mathcal{C}_2 \rightarrow \mathbf{FinStoch}$ defined by the maps

$$\nabla_T = \begin{pmatrix} 0.5 \\ 0.5 \end{pmatrix} \quad \nabla_B = \begin{pmatrix} 0.6 \\ 0.4 \end{pmatrix}$$

$$\boxed{R|TB} = \begin{pmatrix} 0.75 & 0.15 & 0.5 & 0.1 \\ 0.25 & 0.85 & 0.5 & 0.9 \end{pmatrix}$$

We then may compute the causal conditional $[R||T]$ to be

$$Q[R||T] = \begin{array}{c} R \\ | \\ \boxed{R|TB} \\ | \\ \cup \\ \nabla_B \\ | \\ T \end{array} = \begin{pmatrix} 0.51 & 0.34 \\ 0.49 & 0.66 \end{pmatrix}$$

This again agrees with the above assertion that with this causal structure the treatment is effective, as here the probability of recovery with treatment is 0.51, compared with a probability of recovery without treatment of 0.34. In this case the map $Q[R||T]$

is the *only* inferential map from T to R ; it thus may be seen as the only way to deduce information about R from information about T consistent with their causal relationship. Thus within the framework given by the causal theory, there is no possible way to come to the wrong conclusion about the efficacy of the treatment.

In the first case, however, there is one other map; we may infer information about recovery given treatment via

$$\begin{array}{c}
 R \\
 | \\
 \boxed{R|TB} \\
 | \\
 \boxed{B|T} \\
 | \\
 \nabla T \\
 | \\
 T
 \end{array}
 = \begin{pmatrix} 0.51 & 0.34 \\ 0.49 & 0.66 \end{pmatrix}$$

As suggested by the form of the string diagram, this may be interpreted as the chance of recovery if the effect of the treatment on blood pressure is nullified, but nonetheless assuming that the proportion of patients presenting healthy blood pressures at the conclusion of the treatment was typical. In particular, this indicates that if it was inevitable that a group of patients would end up with healthy blood pressure levels in the proportion specified by $[B|T] \circ [T]$, then the treatment would be effective for this group.

Note that the string diagrams themselves encode the flow of causal influence in their depictions of the conditionals. In doing so they make the source of confusion patently clear: we may judge the effect of treatment on recovery in two different ways, one in which we use information about how treatment affected blood pressure, and one in which we forget this link and assume the variables are unrelated.

Finally, observe that causal structures are thus very relevant when interpreting data, and awareness of them can allow one to extract information that could not otherwise be extracted. Indeed, although under the second causal structure the fact that the blood pressure variable biased our selection procedure for treatment—making it more likely that we treated those with unhealthy blood pressure—can be seen as ill-considered experiment design, we see that nonetheless an understanding of the causal structure allowed us to recover the correct conclusion from the data. This becomes critically useful in cases when we do not have the abilities to correct such biases methodologically, such as when data is taken from observational studies rather than controlled experiments.

Chapter 5

The Structure of Stochastic Causal Models

Our aim through this dissertation has been to develop tools to discuss causality, and in particular causal relationships between random variables. Our claim is now these are well described by *stochastic causal models*: models $P : \mathcal{C} \rightarrow \mathbf{CGStoch}$ of a causal theory \mathcal{C} in $\mathbf{CGStoch}$. Indeed, we have seen these are slight generalisations of Bayesian networks in which the factorisation of the joint distribution is made explicit. One advantage of moving to this setting is that we now have a natural notion of map between causal models: a monoidal natural transformation between their functors. We begin this chapter by exploring these, before using the knowledge we gain to look at the existence or otherwise of some basic universal constructions in the category of stochastic causal models.

5.1 Morphisms of stochastic causal models

Fix a causal theory \mathcal{C} . Although we have so far had no problems discussing models of causal theories in \mathbf{Stoch} , we shall define the *stochastic causal models* of \mathcal{C} to be the objects of the category $\mathbf{CGStoch}_{SSM}^{\mathcal{C}}$ of strong symmetric monoidal functors $\mathcal{C} \rightarrow \mathbf{CGStoch}$. This more restrictive definition allows for a more well-behaved notion of maps between stochastic causal models. Indeed, we take the notion of morphism in $\mathbf{CGStoch}_{SSM}^{\mathcal{C}}$ —a monoidal natural transformation between functors—to be the notion of map between stochastic causal models. As we will see in this section, these are much like deterministic stochastic maps. Our aim will be to define the terms in, and then prove, the following theorem:

Theorem 5.1. *Morphisms of stochastic causal models factor into a coarse graining followed by an embedding.*

To this end, let $P, Q : \mathcal{C} \rightarrow \mathbf{CGStoch}$ be stochastic causal models, and let $\alpha : P \Rightarrow Q$ be a monoidal natural transformation. By definition, this means we have a collection of stochastic maps $\alpha_w : Pw \rightarrow Qw$ such that for all variables $w, w' \in \mathcal{C}$ and all morphisms $r : Pw \rightarrow Pw'$ the following diagrams commute:

$$\begin{array}{ccc}
 Pw & \xrightarrow{Pr} & Qw \\
 \alpha_w \downarrow & & \downarrow Qr \\
 Pw' & \xrightarrow{\alpha_{w'}} & Qw'
 \end{array}
 \qquad
 \begin{array}{ccc}
 P\emptyset & \xrightarrow{\alpha_\emptyset} & Q\emptyset \\
 P_* \swarrow & & \searrow Q_* \\
 & * &
 \end{array}$$

$$\begin{array}{ccc}
 Pw \otimes Pw' & \xrightarrow{\alpha_w \otimes \alpha_{w'}} & Qw \otimes Qw' \\
 P_{\otimes, w, w'} \downarrow & & \downarrow Q_{\otimes, w, w'} \\
 Pww' & \xrightarrow{\alpha_{ww'}} & Qww'
 \end{array}$$

We can, however, write this definition a bit more efficiently.

As $P\emptyset$ and $Q\emptyset$ are isomorphic to the monoidal unit, and as the monoidal unit $*$ of $\mathbf{CGStoch}$ is terminal, the above triangle gives no constraints on the morphisms. The lower square specifies the relationships between the maps α_w and $\alpha_{w'}$ and the map $\alpha_{ww'}$ on the product variable. Due to this, it suffices to define the natural transformation α only on the atomic variables v of \mathcal{C} , and let the commutativity of the square specify the maps α_w on the remaining variables. It thus remains to ensure that our maps α_v on the atomic variables v satisfy the defining square of a natural transformation.

We first consider the constraints given by the comonoid maps. The counit maps provide no constraint: since $Q\emptyset$ is terminal, the diagram

$$\begin{array}{ccc}
 Pv & \xrightarrow{\alpha_v} & Qv \\
 \downarrow \bullet & & \downarrow \bullet \\
 P\emptyset & \xrightarrow[\alpha_\emptyset]{\sim} & Q\emptyset
 \end{array}$$

always commutes. On the other hand, the comultiplication maps heavily constrain the α_v : they require that

$$\begin{array}{ccc}
 Pv & \xrightarrow{\alpha_v} & Qv \\
 \downarrow \smile & & \downarrow \smile \\
 Pv \otimes Pv & \xrightarrow{\alpha_v \otimes \alpha_v} & Qv \otimes Qv
 \end{array}$$

or in string diagrams:

The diagram shows an equality between two string diagrams. On the left, a vertical line splits into two, each passing through a box labeled α_v . These two lines then merge back into a single vertical line. On the right, a single vertical line passes through a box labeled α_v . The two diagrams are separated by an equals sign.

The following lemma shows that this is true if and only if each α_v must be deterministic.

Lemma 5.2. *A stochastic map is a comonoid homomorphism if and only if it is deterministic.*

Proof. Let $k : X \rightarrow Y$ be a stochastic map. As the monoidal unit is terminal in **CGStoch**, all stochastic maps preserve the counit. We thus want to show that

The diagram shows an equality between two string diagrams. On the left, a vertical line splits into two, each passing through a box labeled k . These two lines then merge back into a single vertical line. On the right, a single vertical line passes through a box labeled k . The two diagrams are separated by an equals sign.

if and only if k is deterministic.

Now, given $x \in X$ and $B \in \Sigma_Y$, the left hand side of the above equality takes value

$$\begin{aligned} (k \otimes k) \circ \Upsilon(x, B \times B) &= \int k \otimes k(-, -, B \times B) d\Upsilon_x \\ &= \int k(-, B)k(-, B) d\Upsilon_x \\ &= k(x, B)^2, \end{aligned}$$

while the right hand side equals

$$\begin{aligned} \Upsilon \circ k(x, B \times B) &= \int \Upsilon(-, B \times B) dk_x \\ &= \int \chi_B dk_x \\ &= k(x, B). \end{aligned}$$

Thus if k is a comonoid homomorphism, then $k(x, B)^2 = k(x, B)$, and hence $k(x, B) = 0$ or 1 . This shows that k is deterministic. Conversely, if k is deterministic, then $k(x, B)^2 = k(x, B)$ for all $x \in X$, $B \in \Sigma_Y$, so k is a comonoid homomorphism. \square

Summing up, a morphism $\alpha : P \Rightarrow Q$ of stochastic causal models is specified by a collection $\{\alpha_v\}_{v \in V_C}$ of deterministic stochastic maps $\alpha_v : Pv \rightarrow Qv$ such that for all

atomic variables $v \in V_{\mathcal{C}}$ the squares

$$\begin{array}{ccc}
 P(pa(v)) & \xrightarrow{\alpha_{pa(v)}} & Q(pa(v)) \\
 \downarrow P[v|pa(v)] & & \downarrow Q[v|pa(v)] \\
 Pv & \xrightarrow{\alpha_v} & Qv
 \end{array}$$

commute.

We say that a morphism α of stochastic models of \mathcal{C} is an *embedding* if for all objects v of \mathcal{C} the deterministic stochastic map α_v is an embedding. Similarly, we say that a morphism α of stochastic models of \mathcal{C} is a *coarse graining* if for all objects v of \mathcal{C} the deterministic stochastic map α_v is a coarse graining. Theorem 5.1 now follows from Proposition 2.20, with the causal model it factors through having the induced structure.

We caution that despite the similar terminology to deterministic stochastic maps, the situation here differs as stochastic causal models consist of much more data than measurable spaces, and so the compatibility requirements a morphism must obey here are much stricter. For example, while in **CGStoch** it is always possible to find a deterministic map between any two objects, this is rarely possible in **CGStoch** $_{SSM}^{\mathcal{C}}$.

Let $P, Q : \mathcal{C} \rightarrow \mathbf{CGStoch}$ be stochastic causal models, and let $\alpha : P \Rightarrow Q$ be a morphism between them. Then for any prior $[w]$ of \mathcal{C} , the diagram

$$\begin{array}{ccc}
 Pw & \xrightarrow{\alpha_w} & Qw \\
 \uparrow P[w] & & \uparrow Q[w] \\
 * & &
 \end{array}$$

commutes. This says that the pushforward measure of any prior $P[w]$ along the deterministic stochastic map α_w must agree with $Q[w]$. No such map exists, for example, when Pw, Qw are binary discrete measurable spaces and $P[w], Q[w]$ have matrix representations

$$P[w] = \begin{pmatrix} p \\ 1-p \end{pmatrix}, \quad Q[w] = \begin{pmatrix} q \\ 1-q \end{pmatrix}$$

with $p, q \in [0, 1]$ and $q \neq 0, p, 1-p$, or 1.

As diagrams involving all morphisms of \mathcal{C} , and not just the priors, are required to commute, still more constraints apply. Although there are exceptions for finely-tuned parameters, it is generically true that if one wishes to find a coarse graining

$\gamma : P \Rightarrow Q$ between two causal models, then two outcomes of a measurable space Pw can be identified by the map γ_w only when they define the same measure on the codomain for all maps $P\varphi : Pw \rightarrow Pw'$, where $\varphi : w \rightarrow w'$ is any morphism of \mathcal{C} with domain w . The intuition here is that coarse grainings allow us to group outcomes and treat them as a single outcome. But for this to be possible, the outcomes must behave similarly enough to treat as one. Since outcomes now have consequences, we have much higher ability to see differences between them, and hence coarse grainings are far more restrictive for causal models than for measurable spaces.

As embeddings do not identify distinguishable outcomes of the domain, we need not worry about such complications in understanding restrictions on their construction. Nonetheless, if $\epsilon : P \Rightarrow Q$ is an embedding of stochastic causal models, then since the push-forward measure of any prior $P[w]$ along the deterministic stochastic map ϵ_w must agree with $Q[w]$, any measurable set of Qw not intersecting the image of **CGStoch**-embedding must have $Q[w]$ -measure zero. Furthermore, for the naturality squares of the morphisms $\varphi : w \rightarrow w'$ of \mathcal{C} to commute, each map $Q\varphi$ must behave as $P\varphi$ on the image of ϵ_w . This means that an embedding $\epsilon : P \Rightarrow Q$ forces the priors of P and Q to be the ‘same’ up to sets of measure zero.

5.2 Basic constructions in categories of stochastic causal models

In this final section we continue our characterisation of categories of stochastic causal models by exploring a few universal constructions. In particular, we show that these categories have a terminal object, but no initial object, and in general no products or coproducts either. Again fix a causal theory \mathcal{C} .

Proposition 5.3. *The functor $T : \mathcal{C} \rightarrow \mathbf{CGStoch}$ sending all objects of \mathcal{C} to the monoidal unit $*$ of **CGStoch** and all morphisms of \mathcal{C} to the identity map on $*$ is a terminal object in the category $\mathbf{CGStoch}_{SSM}^{\mathcal{C}}$ of stochastic causal models of \mathcal{C} .*

Proof. Note first that, since the monoidal product of $*$ with itself is again $*$, the constant functor $T : \mathcal{C} \rightarrow \mathbf{CGStoch}$ is a well-defined stochastic causal model.

Let $P : \mathcal{C} \rightarrow \mathbf{CGStoch}$ be a stochastic causal model of \mathcal{C} . We construct a monoidal natural transformation $\alpha : P \Rightarrow T$. Then for each $w \in \mathcal{C}$, define $\alpha_w : Pw \rightarrow Tw = *$ to be the unique stochastic map $Pw \rightarrow *$. This exists as $*$ is terminal in **CGStoch**. Furthermore, from the fact that $*$ is terminal in **CGStoch** it is immediate that for each morphism of \mathcal{C} the required naturality square commutes. As these maps

to the terminal object are each deterministic, we thus have a well-defined morphism of stochastic causal models.

By construction it is clear that this is the unique morphism of causal models $P \rightarrow T$. This proves the proposition. \square

The functor T is an example of what we will call a trivial model. Given a measure space (X, Σ, μ) , we define the *trivial model on (X, Σ, μ)* to be the functor $T_\mu : \mathcal{C} \rightarrow \mathbf{CGStoch}$ sending each atomic variable v of \mathcal{C} to X , and each causal mechanism $[v|pa(v)]$ to the map $T_\mu(pa(v)) \rightarrow * \xrightarrow{\mu} Tv = X$ assigning to each element of $T_\mu(pa(v))$ the measure μ . This represents the situation in which all the atomic variables are the same random variable with the same prior, and have no causal influence on each other. We shall use these to show the non-existence of an initial object, products, and coproducts.

Proposition 5.4. *The category $\mathbf{CGStoch}_{SSM}^{\mathcal{C}}$ of stochastic causal models of \mathcal{C} has no initial object.*

Proof. We prove by contradiction. Suppose that $I : \mathcal{C} \rightarrow \mathbf{CGStoch}$ is an initial object of $\mathbf{CGStoch}_{SSM}^{\mathcal{C}}$.

Let $(B, \mathcal{P}(B), \nu)$ be the discrete measure space with two outcomes $\{b_1, b_2\}$ such that the probability of each outcome is one half, and let T_ν be the trivial model of this space. Note that as ν has full support, the only measure space (X, Σ, μ) for which there exists a monic deterministic stochastic map $k : X \rightarrow B$ such that

$$\begin{array}{ccc} X & \xrightarrow{k} & B \\ & \swarrow \mu & \nearrow \nu \\ & * & \end{array}$$

commutes is $(B, \mathcal{P}(B), \nu)$ itself. In this case k must also be an epimorphism in $\mathbf{CGStoch}$; k is either the identity map, or the map $s : B \rightarrow B$ induced by the function sending b_1 to b_2 and b_2 to b_1 . Thus any map of stochastic causal models $\alpha : P \rightarrow T_\nu$ with codomain T_ν must be defined objectwise by coarse grainings of $\mathbf{CGStoch}$, and hence itself be an epimorphism in $\mathbf{CGStoch}_{SSM}^{\mathcal{C}}$.

In particular, the unique morphism of stochastic causal models $\theta_{T_\nu} : I \Rightarrow T_\nu$ must be an epimorphism. Since, by uniqueness, the diagram

$$\begin{array}{ccc}
 T_\nu & \xrightarrow{\alpha} & T_\nu \\
 \theta_{T_\nu} \swarrow & & \searrow \theta_{T_\nu} \\
 & I &
 \end{array}$$

must commute for any morphism of stochastic models $\alpha : T_\nu \Rightarrow T_\nu$, this implies that the only such morphism is the identity map. But it is readily observed that defining $\alpha_v = s : B \rightarrow B$ for each atomic variable v of \mathcal{C} gives a monoidal natural transformation $\alpha : T_\nu \Rightarrow T_\nu$ not equal to the identity. We thus have a contradiction, and so $\mathbf{CGStoch}_{SSM}^{\mathcal{C}}$ has no initial object, as claimed. \square

Example 5.5 (Two objects which have no product). We again work with T_ν , where $(B, \mathcal{P}(B), \nu)$ be the discrete measure space with two outcomes $\{b_1, b_2\}$ such that the probability of each outcome is one half. We will see that the product of T_ν with itself does not exist. Suppose to the contrary that a product X does exist, with projections $\pi_1, \pi_2 : X \Rightarrow T_\nu$. We assume without loss of generality that each Xv is empirical; recall that this means that each point of the set Xv is measurable.

Given the identity monoidal natural transformation $\text{id} : T_\nu \Rightarrow T_\nu$, there exists a unique monoidal natural transformation $\theta : T_\nu \Rightarrow X$ such that

$$\begin{array}{ccc}
 & & T_\nu \\
 & \text{id} \nearrow & \uparrow \pi_1 \\
 T_\nu & \xrightarrow{\theta} & X \\
 & \text{id} \searrow & \downarrow \pi_2 \\
 & & T_\nu
 \end{array}$$

commutes. This shows that for each atomic variable $v \in \mathcal{C}$, Xv has an outcome x_1 of measure $\frac{1}{2}$ such that θ_v is induced by a function mapping b_1 to x_1 , and π_{1v} and π_{2v} are induced by functions mapping x_1 to b_1 . Similarly, we also have $x_2 \in Xv$ of measure $\frac{1}{2}$ such that θ_v is induced by a function mapping b_2 to x_2 , and π_{1v} and π_{2v} are induced by functions mapping x_2 to b_2 . Note that each Xv then has no other outcomes of positive measure.

Let now $\alpha : T_\nu \Rightarrow T_\nu$ be the monoidal natural transformation of the previous proof defined by $\alpha_v = s : B \rightarrow B$, where s is induced by the function sending b_1 to b_2 and

b_2 to b_1 . Then there exists a unique monoidal natural transformation $\theta' : T_\nu \Rightarrow X$ such that

$$\begin{array}{ccc}
 & & T_\nu \\
 & \nearrow \alpha & \uparrow \pi_1 \\
 T_\nu & \xrightarrow{\theta'} & X \\
 & \searrow \text{id} & \downarrow \pi_2 \\
 & & T_\nu
 \end{array}$$

commutes. Now as $\nu(b_1) = \nu(b_2) = \frac{1}{2}$, for each atomic $v \in \mathcal{C}$ the stochastic map $\theta'_v : B \rightarrow Xv$ must then be induced either by the function mapping b_1 to x_1 and b_2 to x_2 , or by the function mapping b_1 to x_2 and b_2 to x_1 . Both cases give a contradiction. In the first case, the composite $(\pi_1 \circ \theta')_v$ is then equal to the identity map on B , contradicting the definition of α . In the second case, the composite $(\pi_2 \circ \theta')_v$ is equal to s , and hence not equal to id_B as required.

Thus no product stochastic causal model $T_\nu \times T_\nu$ exists.

Example 5.6 (Two objects which have no coproduct). Let T be the terminal object of $\mathbf{CGStoch}_{SSM}^{\mathcal{C}}$, and let T_λ be the trivial model on the Lebesgue measure $([0, 1], \mathcal{B}_{[0,1]}, \lambda)$ of the unit interval. We show that these two stochastic causal models have no coproduct in $\mathbf{CGStoch}_{SSM}^{\mathcal{C}}$. To this end, suppose that a coproduct X does exist, with injections $i_1 : T_\lambda \Rightarrow X$ and $i_2 : T_* \Rightarrow X$. We again assume without loss of generality that each Xv is empirical.

To show the difficulties in constructing a coproduct, we use the test object T_μ defined as the trivial model on the measure space $(B, \mathcal{P}(B), \mu)$ with $B = \{b_1, b_2\}$, $\mu(b_1) = 1$ and $\mu(b_2) = 0$. Note that there is a unique map $\beta : T \Rightarrow T_\mu$; this is induced on each atomic $v \in \mathcal{C}$ by the function sending the unique point $*$ of Tv to $b_1 \in B = T_\mu v$. This is the only such map as, since $*$ is a point of measure 1, its image must be a point measure on a point of measure 1. Note also this implies that for each $v \in \mathcal{C}$ the set Xv consists of a point x_1 such that a measurable subset of Xv has $X[v]$ -measure 1 if $x_1 \in Xv$, and measure 0 otherwise.

Consider now maps $\alpha : T_\lambda \Rightarrow T_\mu$. These are defined by, for each atomic $v \in \mathcal{C}$, a choice of a Lebesgue measure 0 subset of $[0, 1]$. We then may let $\alpha_v : [0, 1] \rightarrow B$ be induced by the function mapping each element of this measure zero subset to b_2 , and then remaining elements to b_1 . In particular, for each $p \in [0, 1]$, let $\alpha_p : T_\lambda \Rightarrow T_\mu$ be the monoidal natural transformation such that for all atomic $v \in \mathcal{C}$ the map $(\alpha_p)_v : [0, 1] \rightarrow B$ is induced by the function mapping $p \rightarrow b_2$ and each element of

$[0, 1] \setminus \{p\}$ to b_1 . By the universal property of the coproduct, for each such map there exists a unique map $\theta : X \Rightarrow T_\mu$ such that

$$\begin{array}{ccc}
 T_\lambda & & \\
 \Downarrow i_1 & \searrow \alpha_p & \\
 X & \xrightarrow{\theta} & T_\mu \\
 \Uparrow i_2 & \nearrow \beta & \\
 T & &
 \end{array}$$

commutes. This implies that for each atomic v the function inducing the deterministic stochastic map $i_{1v} : [0, 1] \rightarrow Xv$ does not map p to x_1 . But this implies that the push-forward measure of λ along i_{1v} is the zero measure, contradicting the commutativity of the diagram

$$\begin{array}{ccc}
 [0, 1] & \xrightarrow{i_{1v}} & Xv \\
 & \searrow \lambda & \nearrow X[v] \\
 & * &
 \end{array}$$

This shows that T and T_λ do not have a coproduct in $\mathbf{CGStoch}_{SSM}^c$.

Further Directions

In arriving at this point we have seen that causal theories provide a framework for reasoning about causal relationships, with the morphisms of these categories representing methods of inference, the topology of the string diagrams for these morphisms providing an intuitive visualisation of information flow, and the stochastic models of causal theories slight generalisations of Bayesian networks.

There are many directions in which this study could be continued. One obvious avenue for further exploration is to continue the work of the previous chapter in the characterisation of categories of stochastic causal models. This should, at the very least, provide additional insight into relationships between Bayesian networks. Although we have seen that products and coproducts do not exist in the category of stochastic causal models, and it is likely similar arguments show other types of limits and colimits do not exist, one suggestion is to examine ideas of families and moduli of stochastic causal models. For this, call Bayesian networks equivalent if there are measure-preserving measurable functions between their joint probability distributions that compose to the identity almost everywhere, and call two stochastic causal models equivalent if their induced Bayesian networks are equivalent. It may then be possible to put some geometric structure on the set of stochastic causal models, and subsequently define a moduli problem. This will perhaps generalise work on the algebraic geometry of Bayesian networks, such as that in [8]. One could also explore the relationships between the categories of stochastic causal models of distinct causal theories. Here one might define a functor between such categories if there exists a map of directed graphs between their underlying causal structures.

A weakness of causal theories is that their morphisms only describe predictive inference; reasoning that infers information about causes from their consequences. In general we are interested in other modes of inference too, and extension of the framework to allow discussion of these would make it much more powerful. In the probabilistic case, it can be shown that all conditionals of a joint distribution can be written as morphisms if one can also write Bayesian inverses of the causal conditionals.

Given variables w, w' , these may be characterised as maps $k : w' \rightarrow w$ such that

Methods for constructing such maps often run into issues of uniqueness on outcomes of measure zero in the prior. While in Coecke and Spekkens [4] give a method for realising the Bayesian inverse of a finite stochastic map as transposition with respect to a compactness structure in **Mat** when the prior is of full support, and Abramsky, Blute, and Panangaden [1] give a category, similar to **Stoch**, in which Bayesian inversion may be viewed as a dagger-functor, work remains to be done to merge these ideas with those presented here.

Another topic deserving investigation is suggested by the fact that, although a joint probability distribution is compatible with a causal structure if it satisfies the required set of conditional independence relations, not every possible combination of conditional independence relations of a set of random variables can be represented by a causal structure. Indeed, the number of combinations of conditional independence relations grows exponentially in number of atomic variables, while the number of causal structures grows only quadratically. It is possible that the richer structure of categories may allow us to define causal theories more general than those arising from causal structures, such that models in some category are those that satisfy precisely a given set of conditional independencies, and no more.

Finally, Sections 4.2 and 4.3 suggest their own further lines of investigation. While we have focussed on models in **Stoch** and its subcategories, it would also be worthwhile to understand more thoroughly models in **Rel**, and models in the category **Hilb** of Hilbert spaces and linear maps may be interesting from the perspective of quantum theory. It would also be interesting to find further examples of applications of the graphical languages for causal theories. One option is to look at representations of algorithms used on Bayesian networks, such as Gibbs sampling in Bayesian networks [10].

Bibliography

- [1] Samson Abramsky, Richard Blute, and Prakash Panangaden. Nuclear and trace ideals in tensored $*$ -categories. *Journal of Pure and Applied Algebra*, 2000.
- [2] Samson Abramsky and Bob Coecke. A categorical semantics of quantum protocols. In *Proceedings of 19th IEEE conference on Logic in Computer Science*, pages 415–425. IEEE Press, 2004.
- [3] Robert B. Ash. *Real Analysis and Probability*. Academic Press, 1972.
- [4] Bob Coecke and Robert Spekkens. Picturing classical and quantum Bayesian inference. *Synthese*, 186:651–696, 2012. 10.1007/s11229-011-9917-5.
- [5] Jared Culbertson and Kirk Sturtz. A categorical foundation for Bayesian probability. *Preprint*, 2012.
- [6] Ernst-Erich Doberkat. Kleisli morphisms and randomized congruences for the Giry monad. *Journal of Pure and Applied Algebra*, 211:638–664, 2007.
- [7] Arnold M. Faden. The existence of regular conditional probabilities: necessary and sufficient conditions. *Annals of Probability*, 13(1):288–298, 1985.
- [8] Luis David Garcia, Michael Stillman, and Bernd Sturmfels. Algebraic geometry of Bayesian networks. *Journal of Symbolic Computation*, 39(3-4):331–355, 2005.
- [9] Michèle Giry. A categorical approach to probability theory. In B. Banaschewski, editor, *Categorical Aspects of Topology and Analysis*, volume 915 of *Lecture Notes in Mathematics*, pages 68–85. Springer-Verlag, 1982.
- [10] Tomas Hrycej. Gibbs sampling in Bayesian networks. *Artificial Intelligence*, 46(3):351–363, 1990.
- [11] André Joyal and Ross Street. The geometry of tensor calculus I. *Advances in Mathematics*, 88(1):55–112, 1991.

- [12] Steffen L. Lauritzen. *Lectures on Contingency Tables*. Aalborg University Press, 2nd edition, 1982.
- [13] F. William Lawvere. The category of probabilistic mappings. *Unpublished*, 1962.
- [14] Saunders Mac Lane. *Categories for the Working Mathematician*. Springer-Verlag, 2nd edition, 1998.
- [15] Prakash Panangaden. Probabilistic relations. In C. Baier, M. Huth, M. Kwiatkowska, and M. Ryan, editors, *Proceedings PROBMIV*, pages 59–74, 1998.
- [16] Dusko Pavlovic. Quantum and classical structures in nondeterministic computation. In Peter Bruza, Don Sofge, and Keith van Rijsbergen, editors, *Proceedings of Quantum Interaction*, volume 5954 of *Lecture Notes in Artificial Intelligence*, pages 143–158. Springer-Verlag, 2009.
- [17] Judea Pearl. Bayesian networks: A model of self-activated memory for evidential reasoning. In *Proceedings, Cognitive Science Society*, pages 329–334, 1985.
- [18] Judea Pearl. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann, 1988.
- [19] Judea Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 2nd edition, 2009.
- [20] Peter Selinger. A survey of graphical languages for monoidal categories. In Bob Coecke, editor, *New Structures for Physics*, pages 275–337. Springer-Verlag, 2011.
- [21] Elias M. Stein and Rami Shakarchi. *Real Analysis: Measure Theory, Integration, and Hilbert Spaces*, volume III of *Princeton Lectures in Analysis*. Princeton University Press, 2005.
- [22] Stanislaw Ulam. Zur Masstheorie in der allgemeinen Mengenlehre. *Fundamenta Mathematicae*, 16:140–150, 1930.
- [23] Nicholas Vakhania and Vazha Tarieladze. Regular conditional probabilities and disintegrations. *Bulletin of the Georgian National Academy of Sciences*, 175(2), 2007.

- [24] Thomas Verma and Judea Pearl. Causal networks: Semantics and expressiveness. In *Proceedings of the Fourth Workshop on Uncertainty in Artificial Intelligence*, pages 352–359, 1988.
- [25] M. Wendt. The category of disintegrations. *Cahiers de Topologie et Geometrie Differentielle Categoriqes*, 35:291–308, 1994.
- [26] Jon Williamson. *Bayesian Nets and Causality: Philosophical and Computational Foundations*. Oxford University Press, 2005.