

A Primer on Causal Spaces

A Mathematical Axiomatisation of Causality

Junhyung Park

ETH Zürich

Abstract

This monograph presents the theory of causal spaces in the style of lecture notes, aimed at advanced undergraduate and graduate students in mathematics, as well as mathematically-oriented students in statistics, machine learning, economics, etc. It is also written as a reference for researchers working on the mathematical foundations of causality.

Causal spaces aim to provide an axiomatic framework to study the stochastics of causation. At the heart of the theory of causal spaces is the notion of *interventions*: the study of causality necessitates active interventions in the world, in addition to mere passive observations. This is the key—and essentially, the only—difference of our theory of causality compared to *probability theory*, the mathematical theory of *randomness*.

In Chapter 2, we review the background on probability theory needed to make the monograph self-contained; nonetheless, familiarity with measure-theoretic probability is more or less a prerequisite. No prior exposure to causality is required. Accordingly, beyond the introduction chapter, the main text develops causal spaces without reference to other, more well-known causal theories, such as the structural causal models (SCMs) or the potential outcomes.

The monograph is a significant expansion of the author’s previous works ([Park et al., 2023](#); [Buchholz et al., 2024](#); [Park and Zhou, 2025](#); [Park et al., 2026](#)), both in terms of technical content and exposition. Comments, suggestions corrections are welcome at jun.park@inf.ethz.ch.

Contents

Abstract	ii
1 Introduction	1
1.1 Mathematical axiomatisation	1
1.2 Causality	2
1.3 Probabilities	4
1.4 Mathematisation of causality	6
1.4.1 Potential outcomes	7
1.4.2 Structural causal models	8
1.4.3 Axiomatisation in SCMs	11
1.4.4 Discussion	11
1.5 Causal spaces	12
2 Preliminaries on probability theory	21
2.1 Measurable spaces	21
2.2 Measurable functions	25
2.3 Probability measures and probability spaces	31
2.4 Carathéodory extension theorem and product measures	34
2.5 Integration and expectation	44
2.6 Absolute continuity, Radon-Nikodym and Lebesgue decomposition	54
2.7 Conditioning	58
3 Warm-up: Homogeneous causal spaces	66
3.1 Homogeneous causal spaces	66
3.2 Interventions on homogeneous causal spaces	74
3.3 Causal effects	80
3.3.1 Conditional causal effect	89
3.3.2 Relative causal effect and causal independence	99
3.4 Sources & Identification	107
3.4.1 Conditional sources	110
3.4.2 Relative sources	115
4 Causal spaces	119
4.1 The axioms	119
4.2 Heterogeneity of causation	124
4.3 Preservation	129
4.4 Targeted causal kernels	135

1 Introduction

1.1 Mathematical axiomatisation

A mathematical theory often aims to represent the structure of a concept based on a small number of *axioms*. For example, Dedekind and Peano provided an axiomatic foundation of *arithmetic* (Dedekind, 1888; Peano, 1889), based on which basic operations and concepts on natural numbers can be defined:

- 0 is a natural number.
- For every natural number, there is a successor, which is also a natural number.
- The successors of two different natural numbers are never identical.
- 0 is not the successor of a natural number.
- If a set contains 0 and every successor then it contains every natural number.

Another example is the axiomatic framework of *Euclidean geometry*. In its original form postulated by Euclid in his celebrated book *Elements*, it consisted of the following 5 axioms:

- To draw a straight line from any point to any point.
- To produce (extend) a finite straight line continuously in a straight line.
- To describe a circle with any centre and distance (radius).
- That all right angles are equal to one another.
- That, if a straight line falling on two straight lines makes the interior angles on the same side less than two right angles, the two straight lines, if produced indefinitely, meet on that side on which the angles are less than two right angles (Fig. 1.1).




Figure 1.1: The fifth axiom in Euclidean geometry, the parallel postulate.

This was later modernised by Hilbert (1899). These axioms capture, in a minimal and idealised way, the fundamental geometric notions such as points, lines, circles, and planes. Based on these axioms, the entire theory of Euclidean geometry, including numerous additional definitions and associated results, is built.

As can be seen in the two elementary examples above, a mathematical axiomatisation should clearly capture the essential structural properties of a concept, stripped of accidents, special cases or modelling assumptions. A useful one should then give rise to a coherent and fruitful theory, leading to further natural definitions (that are not primitive objects in the axioms), proving non-trivial theorems, providing canonical examples and enabling conjectures. It should allow the theory to be as broad and general as possible, so that no intended examples are excluded, while being sharp enough to be logically consistent and internally coherent, and to exclude what we do not believe is an instance of the concept being axiomatised. There are a number of modern examples of successful mathematical axiomatisations, such as the axioms of set theory (Jech, 2003), of groups, rings and modules in abstract algebra (Lang, 2012), of vector spaces, inner products, norms, metrics and topologies in analysis (Rudin, 1991), and, most relevantly for us, of probabilities (Kolmogorov, 1933)—see Section 1.3.

A widely accepted mathematical axiomatisation generally comes long after people notice, use, philosophise and even mathematise the concept. It is beyond doubt that early humans were counting objects and doing elementary arithmetic with them, long before the formal axiomatisation of Dedekind and Peano. They similarly thought about various shapes that they saw naturally occurring, such as the (nearly) straight line that a tree makes, or those that they manufactured themselves, such as the triangular tip of an arrow. Hence, it is not that a mathematical axiomatisation is necessary for a concept to be used or to be reasoned with. What it provides is a sharp clarification of the structural essence of an object or a concept, separating the primitive axioms from their logical consequences, and enabling results that apply to all allowed instances, rather than studying each example separately. A successful axiomatisation can create a branch of pure mathematics, or provide a formal foundation for an existing one.

The aim of the present manuscript is to provide a mathematical axiomatisation of the concept of *causality*, which we now describe in detail.

1.2 Causality

Causality is a central pillar of human thought, shaping our actions, decisions and emotions in our daily lives (Woodward, 2005; Waldmann, 2017). The concept is simple and intuitive to the human mind: does A cause B, or more precisely, does a change in A lead to a change in B? For example, does adding salt to my dinner make it saltier? Does daily exercise improve my mood? The essence behind causality is the notion of *interventions*—instead of mere passive observations of the world, we actively intervene on the world, and then observe what subsequently occurs.

As a demonstration of its importance and ubiquitous presence, see how misjudged causation would lead to absurd conclusions. One would see the seat belt sign being turned on every time there is turbulence on an aeroplane, and think, “I wish the pilot would not turn on the seat belt sign. Every time she turns it on, the plane starts shaking!” One may also observe that people who go to see their doctor most often are also those who are the most seriously ill, and conclude “I must not visit my doctor if I

am to avoid being seriously ill”.

It is then undoubtedly a blessing that human cognition is endowed with the ability to reason with causation. Arguably, the fact that we can think causally, act on the basis of causal knowledge, and draw causal conclusions from past events and experiments is a hallmark of human intelligence. Even when the causality in question is not as obvious as in the two examples in the previous paragraph, we strive every day to find causal connections between events, by performing certain actions and seeing what subsequently happens—indeed, causal exploration and discovery form an important part of children’s cognitive development (Muentener and Bonawitz, 2017; Goddu and Gopnik, 2024).

Causality also forms the backbone of virtually all scientific enquiries. The story of Isaac Newton discovering the causal effect of gravity by observing an apple falling from a tree is a famous example, and indeed, the physical laws of the universe encode causal relationships among various objects and quantities (Feynman et al., 1964). Likewise, the mechanisms of chemical reactions (Atkins et al., 2023), the regulation and expression of genes in biology (Alberts et al., 1983), the transmission of traits and diseases in populations (Anderson and May, 1991), the behaviour of organisms in zoology (Alcock, 2009) and the complex interactions within ecosystems in ecology (Begon and Townsend, 2020) all involve questions revolving around causation.

It also plays a crucial role in the functioning of our societies; for example, economists are interested in how the supply and demand of a product causally affect each other (Angrist and Pischke, 2009), lawyers are interested in who caused an event in order to assign blame and responsibility (Hart and Honoré, 1985; Dawid et al., 2017), political scientists are interested in the causal effects of institutions and policies on political behaviour and collective outcomes (Morton and Williams, 2010), epidemiologists are interested in the causes of disease and the effects of treatments and exposures on population health (Rothman et al., 2008), and sociologists are interested in how social structures and environments causally shape human behaviour and life chances (Morgan, 2013). Causality has also been the subject of organised campaigns of doubt with profound societal consequences. For example, tobacco companies spent decades attempting to manufacture uncertainty about the causal link between smoking and lung cancer (Proctor, 2012; Brandt, 2012). Similarly, fossil fuel companies and allied organisations have funded communications and lobbying efforts that cast doubt on, or delay public and policy recognition of, the causal link between fossil-fuel combustion and anthropogenic climate change (Brulle, 2018; Franta, 2021).

Given its centrality and pervasiveness, it would be natural to surmise that efforts to mathematise the concept, and to use it to analyse causality in the world in a quantitative manner, will have a long history. Surprisingly, this is far from the case. Attempts to provide an axiomatisation of it, in the spirit of the previous section, are even rarer still. We will return to these later, but first, we take a digression to the theory of *probabilities*.

1.3 Probabilities

Probability theory is the branch of mathematics that formalises *randomness* (or *uncertainty*, or *stochasticity*) by assigning probabilities to events, and studies the random phenomena and data they generate. For example, a coin toss is a random process that will yield either “Heads” or “Tails”, and we can assign probabilities of $\frac{1}{2}$ to both of those events. The theory allows us to answer, with elegant mathematics, questions such as “how many Heads do we expect to see out of 10 tosses of the coin?” or “how many tosses will we need on average until we see 10 Tails?”

Randomness is evidently another cornerstone of human thought. Unlike causality, this is a concept that enjoys a widely accepted mathematical axiomatisation, based on measure theory (Kolmogorov, 1933). The foundational framework in probability theory is a probability space, namely, a triple $(\Omega, \mathcal{H}, \mathbb{P})$, where the primitive objects are

- Ω is a set of possible outcomes;
- $\mathcal{H}$ is a non-empty collection of events, i.e. subsets of Ω whose probabilities we are interested in measuring; and
- $\mathbb{P}$ is a probability measure on the events.

The axioms that these primitive objects have to satisfy are (see Chapter 2 for more details):

- If $A \in \mathcal{H}$, then $\Omega \setminus A \in \mathcal{H}$; that is, if we are interested in the occurrence of an event, we are also interested in the non-occurrence of the event.
- If $A_1, A_2, \dots \in \mathcal{H}$, then $\cup_{n=1}^{\infty} A_n \in \mathcal{H}$; that is, if we are interested in a countable collection of events, we are also interested in whether or not any of them occurs.
- $\mathbb{P}(A) \in [0, 1]$ for all $A \in \mathcal{H}$, $\mathbb{P}(\emptyset) = 0$ and $\mathbb{P}(\Omega) = 1$; that is, the probability of an event is a value between 0 and 1; the probability of nothing happening is 0 and the probability of something happening is 1.
- For disjoint $A_1, A_2, \dots \in \mathcal{H}$, $\mathbb{P}(\cup_{n=1}^{\infty} A_n) = \sum_{n=1}^{\infty} \mathbb{P}(A_n)$; that is, for events that cannot occur simultaneously, the probability of one of them occurring is the sum of their probabilities.

This axiomatisation did not appear out of the blue. Scholars in the ancient times, such as Aristotle, were already grappling with chance and uncertainty (Franklin, 2015). In the centuries before the work of Kolmogorov, mathematicians had already discovered many of the useful special cases, such as the Gaussian and Poisson distributions, and proved several fundamental results, such as the law of large numbers, limit theorems and the Bayes formula (Hald, 2005, 1998). Even the axiomatisation itself was far from a solo effort—other mathematicians proposed axiomatisations of probabilities in the early 20th century, and the essential measure-theoretic ingredients of Kolmogorov (1933), in particular σ -additivity, trace back to the works of Borel (1909), Henri Lebesgue, Johann

Radon and Maurice Fréchet. For detailed historical accounts leading up to [Kolmogorov \(1933\)](#), we refer the readers to, for example, [Shafer and Vovk \(2018\)](#); [Sheynin \(2018\)](#).

Nor did the axiomatisation replace or invalidate previous works. Rather, it provided a unifying home to known objects and results (in particular, discrete and continuous probabilities were given a common mathematical language), resolved paradoxes (notably, concerning conditional distributions), and enabled a rigorous treatment of complex objects like continuous-time stochastic processes. Above all, it allowed probability theory to be established as a branch of *pure mathematics*. Almost a century later, this axiomatic account of probabilities based on measure theory is considered to be *the* theory of probabilities, recounted in numerous textbooks and taught in classrooms around the world (for modern expositions, we refer the readers to [Çinlar \(2011\)](#); [Billingsley \(1976\)](#); [Klenke \(2008\)](#); [Kallenberg \(1997\)](#); [Durrett \(2019\)](#)).

A closer look shows that these axioms are perfectly aligned with our discussion on the desiderata of mathematical axiomatisations in Section 1.1. They capture the structural essence of probabilities, in the sense that, if these are violated, then we do not believe the mathematical object corresponds to what we know as a probability. Indeed, it is absurd to think of a probability to be negative, or greater than 1, or non-additive over mutually exclusive events. At the same time, the axioms remain minimal and general, so they do not exclude objects that we would still wish to call probabilities. The influence of the theory now extends far beyond pure mathematics, spanning many domains that were surely beyond the imagination of its original developers.

Returning to our main subject, probability theory alone has no innate mechanism for encoding causal information. A classic example of the mantra “correlation does not imply causation” is that ice cream sales and the number of shark attack incidents are highly correlated (see Fig. 1.2), but clearly, there is no causation between them. In other words, forcing people to buy more ice cream in an arbitrary month will lead to no changes in the number of shark attacks (recall, the essence of causality is what happens under *interventions*), and vice versa. A probability measure $\mathbb{P}$ depicting this situation would reflect a high level of dependence between the ice cream sales and shark attack variables (see Section 2.3 for a precise definition of (in)dependence), but does not have the capacity to encode the fact that there is no causation between the two. We stress that this is not a defect of the theory, as it was never intended for that purpose.

Nevertheless, probability remains indispensable to most theories of causality, including our own, because causal relationships are rarely deterministic. Smoking, for example, “causes” lung cancer in the sense that it increases its probability, but it does not make it




Figure 1.2: In the hotter months, people are more likely to buy ice cream, as well as to go on holidays to the sea, leading to high correlation.




Figure 1.3: Sewall Wright’s path diagram, illustrating causal relationships between litter mates and their parents, with variables representing genetic constitutions, environmental factors and other ontogenetic irregularities (Wright, 1920).

inevitable. A probabilistic description can capture the fact that smoking and lung cancer are associated, but it cannot by itself encode the crucial asymmetry between smoking causing lung cancer and lung cancer causing smoking. In other words, probability theory can describe patterns of co-occurrence, but not causal direction (if one exists; otherwise, the lack thereof). Thus, if causality is to be treated mathematically, a theory genuinely extending probability is required.

1.4 Mathematisation of causality

As alluded to at the end of Section 1.2, causality does not have such a rich history as probability theory—in fact, it was often dismissed as a serious scholarly subject. Russell (1912) famously wrote in the opening sentence: “I wish, first, to maintain that the word ‘cause’ is so inextricably bound up with misleading associations as to make its complete extrusion from the philosophical vocabulary desirable”. Judea Pearl, the Turing award winner for his work on causality (more on him imminently) also describes the following anecdote in his popular science book (Pearl and Mackenzie, 2018): “Even two decades ago, asking a statistician a question like ‘Was it the aspirin that stopped my headache?’ would have been like asking if he believed in voodoo. To quote an esteemed colleague of mine, it would be ‘more of a cocktail conversation topic than a scientific inquiry.’” Hence, in spite of works by some early pioneers (notably, Wright (1920) in biology—see Fig. 1.3—and Haavelmo (1943) in econometrics), the endeavour of mathematising the concept of causality only began to take its modern and sustained form towards the end of the 20th century, with the introduction of Donald Rubin’s potential outcomes (Rubin, 1974, 1978; Holland, 1986) (see (Imbens and Rubin, 2015; Hernán and Robins, 2020)

for recent expositions) and Judea Pearl’s structural causal models (Pearl, 1988, 1995) (see (Pearl, 2009; Peters et al., 2017) for recent expositions). Let us take a closer (but informal) look at each of them.

1.4.1 Potential outcomes

The potential outcomes framework of Donald Rubin and his collaborators usually starts with a treatment variable W , often binary (for example, indicating whether or not a patient receives a medicine), and a collection of *potential outcome* variables Y_w indexed by possible values w of W (for example, the patient’s blood pressure after treatment w). The causal interpretation here is that Y_w is the outcome value that *would* be observed *if* the intervention $W = w$ were to take place (we stress again that interventions are the crux of causality). We also have the *observed* outcome variable Y , with the near-universal *consistency* assumption that

$$Y = \sum_w \mathbf{1}(W = w) Y_w,$$

where the sum ranges over the possible values of the treatment variable W . Further, we have an optional random vector X of *covariates* (for example, representing the patient’s characteristics such as weight, gender, etc.). The framework is almost always explicit about the i.i.d. copies of these variables, $W(i), Y_w(i), Y(i), X(i)$ for $i = 1, \dots, n$, representing the “units” (patients) under consideration. Then, the prototypical quantities of interest are the *average treatment effect*, defined (in the binary treatment case) as

$$\text{ATE} = \mathbb{E}[Y_1 - Y_0]$$

and the conditional average treatment effect,

$$\text{CATE}(X) = \mathbb{E}[Y_1 - Y_0 \mid X],$$

for when one wishes to take into account the heterogeneity of the treatment effect depending on the covariates.

It is possible to consider essentially arbitrary potential outcome variables (Holland, 1986; Ibeling and Icard, 2023; Mlodozieniec et al., 2025), corresponding to interventions on any variables in the system, rather than having a fixed treatment variable. Then, simply by placing a joint probability measure over this collection, a very high level of flexibility can be guaranteed. This is powerful, but it also means that the formalism encodes a joint distribution over responses to different interventions even in settings where only individual intervention-specific distributions, or simple comparisons between them, are substantively of interest.

By this point, the reader will have realised that the potential outcomes framework is not, strictly speaking, a new foundational mathematical framework. It remains entirely within the realm of probability theory, while assigning causal and counterfactual *interpretations* to certain random variables. The issue is not that the potential outcomes framework is formulated within probability theory; any modern axiomatic theory will

rest on prior mathematics. Rather, its causal content lies in the interpretation of certain random variables as arising from interventions, not in a collection of primitive causal objects and axioms whose role is to isolate causal structure itself.

An analogue would be to develop probability theory using only real numbers. In simple examples, one may interpret the number $\frac{1}{2}$ as the probability of “Heads” in a coin toss and reason successfully about randomness, even without first introducing the measure-theoretic axioms of probability. It is undoubtedly valuable that one is able to do this—it would be absurd to have to learn measure theory before being able to reason with randomness and probabilities at all. The potential-outcomes framework is similar in spirit: it allows one to reason mathematically about causality by endowing probabilistic objects with causal meaning, without first isolating and axiomatising a mathematical structure whose primitive purpose is to encode causality itself. It is an extremely useful framework, and has proved highly influential, but it is not an axiomatisation of causality in the foundational sense. No axioms are given to capture the structural essence of what an intervention is, or what causation is.

1.4.2 Structural causal models

The structural causal models (SCMs) of Judea Pearl and his collaborators take a different approach. Here, the primitive objects are the exogenous variables $U_1, \dots, U_m$ and endogenous variables $V_1, \dots, V_n$, a probability distribution $\mathbb{P}$ over the exogenous variables, and functions f_i , for $i = 1, \dots, n$, called *structural equations*, corresponding to each endogenous variable, which take a subset of the exogenous and other endogenous variables as arguments. Then, the observational distribution on the endogenous variables is calculated as a pushforward measure of the exogenous distributions through the structural equations, and the interventional distributions are calculated in the same way but with some of the structural equations altered (or *intervened on*).

The SCM framework thus encodes the causal information directly in the structural equations. The interventions of interest do not have to be encoded a priori, unlike the potential outcomes framework. One often associates a graphical representation to an SCM, with the vertices being the endogenous variables and a directed edge being present from V_i to V_j if V_i enters as an argument in f_j . The graph provides a visual, intuitive understanding of the data generating process and the causal relationships (see Fig. 1.4 for an example). A particular feature of the SCM framework is that it comes with graphical criteria for *identification*, i.e. inferring interventional distributions from the observational distribution.




Figure 1.4: An example of the graphical representation of an SCM. Each directed edge represents a causation from the parent node to the child, i.e. changing the distribution of the parent affects the distribution of the child.

While an intuitive, flexible and useful framework, SCMs nevertheless fall short of a foundational axiomatic theory in the sense pursued in this monograph, not only because there are no axioms that attempt to capture the structural essence of what an intervention is, but also because it excludes many interesting cases from being realised. In other words, plausible collections of observational and interventional distributions on the endogenous variables extend beyond what can be realised as a pushforward through the structural equations in the above manner.




Figure 1.5: Cycle.

A prototypical example is when there exist cyclic causal relationships, which are very common. For example, in economics, the demand of a product causally affects its supply, and vice versa, and in ecology, prey and predator populations causally affect each other's numbers. When cycles exist, the pushforward mechanism may not ensure the existence or the uniqueness of observational or interventional distributions on the endogenous

variables. An extremely simple configuration of two integer-valued variables X and Y with the following observational and interventional distributions do not have an SCM representation:

- The pair $(X, Y) = (1, 1)$ has probability 1 in the observational state.
- When X is intervened to have value n , the pair (X, Y) takes value $(n, n + 1)$ with probability 1.
- When Y is intervened to have value n , the pair (X, Y) takes value $(n + 1, n)$ with probability 1.

The issue is that SCMs force the observational and interventional measures both to be propagated through the same structural equations, except those that are intervened upon. This means that cycles are not naturally or uniformly handled in the basic SCM formalism without extra assumptions, and one has to check for solvability, i.e. that the given primitive objects (exogenous distribution and structural equations) give rise to a unique set of observational and interventional probabilities (Bongers et al., 2021; Forré and Mooij, 2020).

For another example of combination of observational and interventional distributions that SCMs cannot represent, consider the following. Suppose $\mathbb{P}(X = 0, Y = 0) > 0$ observationally, but under the intervention to set $X = 0$, one has $\mathbb{P}(Y = 0) = 0$. No SCM can realise this, because any exogenous state yielding $X = 0, Y = 0$ observationally remains a solution after replacing the structural equation for X by the constant equation $X = 0$, forcing $\mathbb{P}^{do(X=0)}(Y = 0) \geq \mathbb{P}(X = 0, Y = 0) > 0$.

Further, the SCM framework is inherently geared towards a finite number of variables, requiring parents as arguments for the structural equations. This renders the realisation of continuous-time stochastic processes cumbersome. Researchers often resort to an additional layer of modelling, for example via differential equations and equilibrium constructions (Mooij et al., 2013; Rubenstein et al., 2018; Bongers et al., 2018; Peters et al., 2022; Lorch et al., 2024; Boeken and Mooij, 2024).




Figure 1.6: Temperature (T), economy (E) and world (W) acting as confounding variables between ice cream sales (I) and the number of accidents (A).

Even when a representation is possible, an SCM quietly inserts unnecessary assumptions that should be avoided in a fully general framework. To illustrate this subtle point, let us return to the example introduced at the end of Section 1.3, and let I and A be random variables representing ice cream sales and the number of accidents respectively. We cannot represent this situation with just the two variables of interest (as in Fig. 1.6(a)), since drawing an arrow between them would imply causation, and not drawing an arrow would imply independence.

One would typically explicitly include the confounding variable, in this case the temperature (T), that causally affects both I and A , as in Fig. 1.6(b), and let the corresponding structural equations f_I and f_A take T as an argument. This creates dependence between I and A without causation as desired, and is perfectly adequate for *modelling* purposes. However, this tacitly bakes in the assumption that there is nothing else in the world other than T that causally affects both I and A , which is likely not true since there are other variables, such as the state of the economy (E), that may also causally affect both I and A (see Fig. 1.6(c)). An axiomatic framework of causality should allow the representation of causal information, such as correlation but no causation in this case, without relying on additional assumptions about the world.

A possible solution could be to include a variable W (for “world”) that collects everything in the world that causally affects both I and A , as in Fig. 1.6(d). Descriptively, this is unobjectionable. However, the numerical value of such a variable, as well as its probability distribution, would have no meaning, and the structural equations f_I and f_A that take W as its argument would also entail no meaning beyond being a construct that gives rise to the desired distributions. The price of expressing, within the SCM framework, the modest claim that I and A are dependent while neither causes the other is thus the specification of an object that is largely arbitrary. An axiomatic framework of causality should be able to carry such causal information directly, without passing through a generating mechanism that the information itself does not determine.

1.4.3 Axiomatisation in SCMs

There are a number of papers, in the tradition of Judea Pearl and Joseph Halpern, which state an axiomatisation as their goal within the SCM framework. However, the flavour of these works is fundamentally different in kind from the goal pursued in this monograph. Let us discuss this distinction in depth, as axiomatisation is a central aspect of our theory.¹

In mathematics, “axioms” refer to the *primitive, defining properties* of a mathematical structure. The group axioms (closure, identity, associativity and inverse) define *what a group is*. The probability axioms as seen in Section 1.3 define *what a probability space is*. This is the kind of axiomatisation we will pursue in this monograph. On the other hand, in logic, the word “axiomatisation” is used to mean a formal proof system for reasoning about statements relative to a fixed class of models. In other words, axioms are not used to define a structure, but rather, they are used to characterise valid reasoning within existing structures. To emphasise the distinction once more: in the former case, axioms define the class of models; in the latter, the class of models is fixed in advance, and axioms characterise the valid formulae over that class.

The papers within the SCM framework that study axiomatisations sit firmly in the second camp (Galles and Pearl, 1997, 1998; Halpern, 2000; Bonet, 2001; Halpern, 2013). In these papers, one *starts* with a class of structural causal models, already equipped with interventional semantics. What are then called “axioms” there are deductive principles governing reasoning within this pre-specified SCM formalism (although the effectiveness axiom is very similar to one of our axioms—see Section 1.5). The soundness and completeness results are proved to hold for fixed classes of SCMs and axioms, showing that the deductive system exactly captures the formulas that are valid in a given class of SCMs. Ibeling and Icard (2020); Beckers et al. (2023); Beckers (2025); Peters and Halpern (2025); Barbero and Sandu (2021); Barbero and Yang (2022); Barbero et al. (2023); Barbero and Virtema (2023); Barbero (2024) largely follow the tradition of earlier papers, but with their own variants or extensions of SCMs.

Singal and Michailidis (2024) also uses the word “axioms” within the SCM framework, but here, the axioms are for the propagation of causal effects, rather than for the primitive objects capturing intervention itself. The approach of Sadeghi and Soo (2025) differs significantly from the above papers, and does not use an SCM as the underlying model, but instead a family of interventional distributions (one interventional distribution per variable). The axioms presented include transitivity and observability, but have the flavour of being *properties* that interventional distributions *may* satisfy to facilitate certain kinds of causal reasoning and drawing of graphs, rather than primitive assumptions defining the underlying objects.

1.4.4 Discussion

The potential outcomes and SCM frameworks have been hugely influential. In numerous domains, one wishes to infer causal effects, or to discover the existence and direction of

¹The author is indebted to Thomas Icard for highly stimulating discussions on this subject.

causal relationships, from data, and the two frameworks have been deployed extensively by statisticians, econometricians, machine learning researchers and many others for precisely these purposes. They are particularly well-suited for the inference of causation from observational data (Rosenbaum, 2002, 2020), and there has always been a heavy focus on identifying causality from observational data. Indeed, a large part of the debate between these two frameworks has revolved around the naturalness and verifiability of the assumptions natively encoded in the respective frameworks that enable such identification (Pearl, 2009; Imbens, 2020; Weinberger, 2023), though works on comparing more fundamental aspects also exist (Markus, 2021; Ibeling and Icard, 2023).

At the same time, causality has not attained the same standing as probability theory as a branch of pure mathematics. In the view of the author, this is due, at least in part, to the absence of a widely accepted object-level mathematical axiomatisation of causality itself, analogous to that of probability theory. Existing frameworks are powerful and fruitful, but they do not isolate primitive mathematical objects whose sole purpose is to encode interventions and their consequences in the way that measure-theoretic probability isolates the mathematical structure of randomness. Our aim in this monograph is to make a first step in that direction by proposing such an axiomatisation. We explain this programme in the next section.

Before we move on, we must be careful not to overstate the envisaged practical significance of our axiomatisation. We expect that its direct impact on practitioners, who carry out the valuable work of inferring and discovering causality from data, will be negligible, while SCMs and potential outcomes will continue to be used extensively. Here, again, a direct comparison can be made with the use of measure-theoretic probability theory by practitioners inferring non-causal information from data, which is very little. However, probability theory and statistics can be seen as two sides of the same coin, with the advances in one field aiding the development of the other over the past century, and our hope is that our axiomatisation can play the same role in the field of causality.

1.5 Causal spaces

Let us repeat our goal, for emphasis. We want to build an axiomatisation of causality, which is fundamentally a study of interventions, and we want to do it in such a way that the randomness after an intervention is captured. For the latter purpose, the basic objects of probability theory—the outcome set Ω , σ -algebra $\mathcal{H}$ of events, and the probability measure $\mathbb{P}$ —are the right ingredients. On the other hand, as we have already established in Section 1.3, this triple $(\Omega, \mathcal{H}, \mathbb{P})$ is not designed to accommodate interventions, so our plan is to equip a probability space with further structure, in order to isolate and capture the essence of interventions. Specifically, the additional structure needs to encode

- (a) where an intervention will take place;
- (b) the effect of an intervention.

For (a), we will require that the measurable space $(\Omega, \mathcal{H})$ is in the form of a product, $(\Omega, \mathcal{H}) = \otimes_{t \in \mathbb{T}} (\Omega_t, \mathcal{H}_t)$, for some index set $\mathbb{T}$. Interventions will then take place on sub- σ -algebras of $\mathcal{H}$ corresponding to subsets of $\mathbb{T}$.

For (b), we need an additional set of primitive objects that capture the effect of interventions. We will deploy transition probability kernels $K_{\mathbb{S}} : \Omega \times \Omega \times \mathcal{H} \rightarrow [0, 1]$ (see Section 2.3 for a precise definition), called *causal kernels*, indexed by subsets $\mathbb{S} \subseteq \mathbb{T}$. For each pair $(\eta, \omega) \in \Omega \times \Omega$, the map $A \mapsto K_{\mathbb{S}}((\eta, \omega); A) : \mathcal{H} \rightarrow [0, 1]$ is a probability measure, to be interpreted as: *starting from the pre-intervention state η , the probability of A after the $\mathbb{S}$ -coordinates are set to $\omega_{\mathbb{S}}$* . For example, $\mathbb{S}$ could represent some medicine, ω the dose of the medicine, η some pre-intervention characteristics of a patient and A an event related to an outcome of interest, such as the patient's blood pressure.

We denote the collection of all causal kernels by $\mathbb{K} = \{K_{\mathbb{S}} : \mathbb{S} \subseteq \mathbb{T}\}$, and call it the *causal mechanism*. Finally, we call the quadruple $(\Omega, \mathcal{H}, \mathbb{P}, \mathbb{K})$ a *causal space*. The fact that interventional distributions are the primitive objects stands in marked contrast from the SCM framework, in which they are secondary objects that are derived from the primitives of structural equations and exogenous distributions. In this regard, causal spaces bear more similarities to the potential outcomes framework, or the decision-theoretic framework of Philip Dawid (Dawid, 2015, 2021). However, neither considers the use of transition probability kernels, which is a distinctive feature of causal spaces that allows us to elegantly treat soft, heterogeneous and targeted interventions within the same framework (see Chapter 5).

There is a subtle design choice that is implicitly made for the causal kernels: that, when intervening on the $\mathbb{S}$ -coordinates for a particular pre-intervention value η , the value ω with which we intervene is all that matters. This might make some readers wonder: surely the *way* in which that intervention value is imposed also matters? For example, suppose we are interested in the effect of a particular drug on a patient. Then giving the patient a placebo and giving the patient nothing sets the drug variable to the same value, zero. However, the effect on the patient will be different. In this work (and most other works on causality), we work on the premise that, if the way in which an intervention takes place matters for the effect of an intervention, then it should be explicitly included in the model, and that, for the coordinates modelled, it is indeed only the value that matters.²

What axioms should the causal kernels satisfy, to capture the essence of an intervention? We want the axioms to impose just enough constraints so that, if those are violated, then we no longer believe that the concept of an intervention is captured, while being as general as possible so as not to exclude any instance of what we believe to be an intervention. The following two axioms are what we present as the minimal requirements for causal kernels intended to represent interventions.

- (i) $K_{\emptyset}((\eta, \omega); A) = \mathbf{1}_A(\eta)$ for all $A \in \mathcal{H}$; that is, if we intervene on no coordinates, then the pre-intervention outcome η remains intact.

²The author thanks Joris Mooij and Alexander Reisach for this discussion. For an interesting way of treating this problem with bipartite graphs, see Joris Mooij's presentation (Mooij, 2026).

- (ii) $K_{\mathbb{S}}((\eta, \omega); A) = \mathbf{1}_A(\omega)$ for all $A \in \mathcal{H}_{\mathbb{S}}$; that is, if we fix the $\mathbb{S}$ -coordinates to the value ω , then the $\mathbb{S}$ -coordinates have value ω .

These axioms are specific to causal kernels, not to transition probability kernels in general. The first says that if nothing is intervened on, nothing changes; the second says that when the $\mathbb{S}$ -coordinates are given values ω during an intervention, then those are the values that the $\mathbb{S}$ -coordinates have after the intervention. We impose no further axioms at the foundational level. These two axioms are uncontroversial—if we want the causal kernels $K_{\mathbb{S}}$ to represent interventions, then these are properties that we would minimally require. We hope to convince the reader, with the content of this monograph, that a coherent, consistent, rich and fruitful mathematical theory emerges based on these minimal axioms, with clear causal semantics.

It must be emphasised that our aim in proposing an axiomatisation of causality is neither to replace nor to criticise the existing frameworks. Their value for the purposes they were designed to address is beyond doubt. Our point is rather that they do not constitute, nor were they intended to constitute, a foundational axiomatic theory of causality. In this sense, causality has not yet come to be regarded as a branch of pure mathematics in the way that probability theory has. The digression on the historical development of probability theory in Section 1.3 was included partly to clarify this analogy. We expect the axiomatic theory developed in this monograph to recover, unify, and generalise many known results from the existing frameworks, much as Kolmogorov’s axiomatisation unified earlier work in probability. At the same time, those frameworks should continue to play an important role in the specification of causal spaces, just as classical distributions such as the Gaussian and Poisson remain fundamental examples within modern probability theory. We also expect that the axiomatic approach will enable a rigorous and elegant treatment of more complex systems, such as continuous-time stochastic processes, for which the existing frameworks are less naturally suited.

A note on counterfactuals Counterfactual thinking refers to the process of simultaneously imagining two or more parallel worlds: statements of the form “If I had studied more, would I have done better in the exam?”, “But for Alec Radford, artificial intelligence would not be in its current state today” are examples of counterfactual statements, as they involve an imaginary world which did not in fact occur. The tradition of possible worlds semantics of counterfactuals has a long history in philosophy (Goodman, 1947; Lewis, 1973, 1986; Stalnaker, 2003), while psychologists have studied how imagining counterfactual scenarios influences emotions, intentions, decisions and moral judgments (Byrne and McEleney, 2000; Epstude and Roese, 2008; Buchsbaum et al., 2012; Van Hoeck et al., 2015; Byrne, 2016; Gerstenberg, 2024). The relationship between causality and counterfactuals has been studied vigorously by philosophers and statisticians alike (Collins et al., 2004; Pearl, 2000): when considering the causal effect of an action, one often compares, implicitly or explicitly, the ensuing events against those in an “imagined” world in which the original action is different.

Consequently, it features prominently in the two mathematical frameworks introduced in Section 1.4. In the potential outcomes framework, the potential outcome variables




Figure 1.7: Left: Pearl’s ladder of causation. Concepts in the upper rungs are strict generalisations of those in the lower rungs. SCMs are used to calculate the observational, interventional and counterfactual distributions in all of the rungs. Right: Causal spaces and counterfactual probability spaces are each orthogonal extensions of probability spaces, and combining the two, we obtain counterfactual causal spaces.

Y_w are each considered to live in separate counterfactual worlds, depending on the value w of the treatment variable; thus, counterfactual semantics are directly built into how causality is interpreted in this framework. On the other hand, the type of counterfactuals typically considered in the usual SCM framework is conditioning in the factual world and intervening in the counterfactual world through the “abduction–action–prediction” procedure (Pearl, 2009). Even those variants that explicitly only treat non-interventional counterfactuals (for example, the extended SCMs of Lucas and Kemp (2015) and back-tracking SCMs of von K  gelgen et al. (2023)), base their mathematics on causal models designed for interventions.

The essence behind counterfactuals is not that we have multiple worlds based on different interventions or conditioning information, but that we have a *joint* distribution over these worlds. Our approach deviates from the existing literature on counterfactuals, in that we do not view causality or interventions as a necessary ingredient to study counterfactuals. Rather, we believe that expanding the outcome set and the σ -algebra of events as a product of factual and counterfactual worlds is the most faithful way to capture counterfactual thinking within the axioms of probability and causal spaces (see Fig. 1.7, right-hand side), orthogonally to the concept of interventions. In particular, counterfactuals do not inherently contain any causal information, which can only be gained from interventions. We call such probability and causal spaces with expanded outcome sets and σ -algebras *counterfactual probability spaces* and *counterfactual causal spaces*, to be introduced formally in the final chapter, Chapter 8.

In a series of papers, Dawid (1999, 2000, 2007) argues that counterfactual considerations are unnecessary and potentially misleading for effects of causes, and only interven-

tional considerations are required (for which he proposes a decision-theoretic framework (Dawid, 2012, 2015, 2021)). We agree that counterfactuals and interventions are orthogonal concepts (as argued above and illustrated in Fig. 1.7), but he goes on to argue that, since counterfactual statements are by nature impossible to corroborate with data, there is no point in developing a theory of it. However, we do not agree that this provides grounds for an outright rejection of a formalism for counterfactuals; even though empirical verification may be impossible, axiomatising such a fundamental component of human thought is still, for reasons we discuss, a worthwhile pursuit.

Index of notation

Below, we collect the notation used throughout this monograph.

- General**
- $\mathbb{N} = \{1, 2, \dots\}$, the set of natural numbers.
 - $i, j, k, n, m \in \mathbb{N}$, natural numbers.
 - $\mathbb{R} = (-\infty, \infty)$, the set of real numbers.
 - $\mathbb{R}_+ = [0, \infty)$, the set of non-negative real numbers.
 - $\overline{\mathbb{R}} = [-\infty, \infty]$, the set of real numbers extended to include $-\infty$ and ∞ .
 - $\overline{\mathbb{R}}_+ = [0, \infty]$, the set of non-negative real numbers extended to include ∞ .
 - $a, b, c, \varepsilon, \alpha$ real numbers.
 - q rational number.

- General set-related notation**
- $\cup$, union.

- $\cap$, intersection.
- $\setminus$, difference.
- $\triangle$, symmetric difference.
- $\mathcal{P}(\cdot)$, power set.
- $\times$ Cartesian product.

- Index sets**
- $\mathbb{T}$, the full index set.
 - $\mathbb{S}$, typically used as the running index (e.g. for all $\mathbb{S} \in \mathcal{P}(\mathbb{T}), \dots$)
 - $\mathbb{U} \in \mathcal{P}(\mathbb{T})$, typically used for the first intervention.
 - $\mathbb{V} \in \mathcal{P}(\mathbb{T})$, typically used for a subsequent intervention.
 - $\mathbb{W} \in \mathcal{P}(\mathbb{T})$, typically used as auxiliary coordinates in relative preservation and relative sources.
 - $\mathbb{I}, \mathbb{J}$, for other uses of index sets.
 - $\chi_t : \mathcal{P}(\mathbb{T}) \rightarrow \{0, 1\}$, the membership indicator function.

- Outcome sets**
- $\Omega = \otimes_{t \in \mathbb{T}} \Omega_t$, the entire outcome set, a product with index set $\mathbb{T}$.
 - $\Omega_{\mathbb{S}}$ for $\mathbb{S} \in \mathcal{P}(\mathbb{T})$, the coordinates of Ω corresponding to $\mathbb{S}$.

- Outcomes**
- $\omega \in \Omega$, typically the outcome at which an intervention takes place.
 - $\omega_t \in \Omega_t$, the coordinate of ω corresponding to $t \in \mathbb{T}$.
 - $\omega_{\mathbb{S}} = \pi_{\mathbb{S}}(\omega) \in \Omega_{\mathbb{S}}$, the coordinates of ω corresponding to $\mathbb{S}$.
 - $\eta \in \Omega$, typically used as the pre-intervention heterogeneity argument.
 - $\gamma \in \Omega$, typically the outcome at which a conditioning takes place.

- σ -algebras**
- $\otimes$ product of σ -algebras.

- $\mathcal{E}_t$, coordinate σ -algebra of Ω_t for each $t \in \mathbb{T}$.
- $\mathcal{H} = \otimes_{t \in \mathbb{T}} \mathcal{E}_t$, full product σ -algebra on Ω .
- $\mathcal{H}_{\mathbb{S}}$, the sub- σ -algebra of $\mathcal{H}$ corresponding to the coordinates $\mathbb{S}$.
- $\mathcal{E}_{\mathbb{S}}$, the marginal σ -algebra on $\Omega_{\mathbb{S}}$.
- $\mathcal{F} \subseteq \mathcal{H}$, typically used as a “causally affected” σ -algebra.
- $\mathcal{G} \subseteq \mathcal{H}$, typically used as a “conditioning σ -algebra”.
- $\mathcal{J} \subseteq \mathcal{H}$, typically used as heterogeneity σ -algebras.
- $\sigma(\cdot)$, the generated σ -algebra.
- $\vee$, the σ -algebra generated by the union.
- $\mathcal{B}$, Borel σ -algebra on $\mathbb{R}$, $\overline{\mathbb{R}}$, $\mathbb{R}_+$ and $\overline{\mathbb{R}}_+$.
- $\mathcal{T}$, the σ -algebra on $\mathcal{P}(\mathbb{T})$ generated by the collection of functions $\{\chi_t : t \in \mathbb{T}\}$.
- $\mathcal{U} = (\mathcal{T} \otimes \mathcal{H} \otimes \{\emptyset, \Omega\}) \vee \sigma(\{(\mathbb{S}, \eta, \omega) : t \in \mathbb{S}, \omega \in H\} : t \in \mathbb{T}, H \in \mathcal{H}_t)$.
- $\Sigma_{\mathbb{U}}(A)$ and $\Sigma_{\mathbb{U}}$, the heterogeneity families of $K_{\mathbb{U}}$.
- τ_U , the targeting scope of U .
- $\tau_{U,J}$, the assignment scope of J .
- $\mathcal{H}|_A$ for $A \in \mathcal{H}$, the trace of $\mathcal{H}$ on A .

Other collections of subsets • $\mathcal{C}, \mathcal{D}, \mathcal{M} \subseteq \Omega$, often π -systems or λ -systems.

- $\mathcal{S}$, semi-rings.
- $\mathcal{R}$, algebras.

Measurable Sets • $A \in \mathcal{H}$, typically used as a “causally affected” event.

- $B \in \mathcal{H}$, typically in $\mathcal{H}_{\mathbb{U}}$ where the intervention takes place.
- $C \in \mathcal{H}$, typically used as a collection of $\omega \in \Omega$ whence causation issues.
- $G \in \mathcal{H}$, typically used as a “conditioning event”.
- $H \in \mathcal{H}$, typically used as a set in the conditioning σ -algebra.
- $D, E, F \in \mathcal{H}$, for other uses of measurable sets.
- $\Lambda^1, \Lambda^2, \Lambda^3 \subseteq \Omega$, typically used for Lebesgue decompositions.
- A_{ω} , sections.

Other sets • $S, T \in \mathcal{S}$, sets in semi-rings.

- $R \in \mathcal{R}$, sets in algebras.

Measures • $\mathbb{P}$, the observational measure on $(\Omega, \mathcal{H})$.

- $\mathbb{E}$, expectation with respect to $\mathbb{P}$.
- $\mathbb{Q}$, typically an assignment measure on $(\Omega, \mathcal{H}_{\mathbb{U}})$.

- δ_ω , the Dirac delta measure at the outcome ω .
- μ, ν , other (possibly non-probability) measures or set functions.

Functions

- f, g, h functions, often measurable.
- $\mathcal{F}$, class of functions.
- $\text{Im}(f) = \{f(\omega) : \omega \in \Omega_1\}$, image of $f : \Omega_1 \rightarrow \Omega_2$.
- $f^{-1}(A) = \{\omega \in \Omega : f(\omega) \in A\}$, pre-image of a set $A \subseteq \Omega_2$ under f .
- $\pi_{\mathbb{S}} : \Omega \rightarrow \Omega_{\mathbb{S}}$, the canonical projection.
- $\iota_{\omega_{\mathbb{S}}} : \Omega_{\mathbb{T} \setminus \mathbb{S}} \rightarrow \Omega$, insertion.
- $\mathbf{1}_A$, the indicator function of the measurable set A .
- $\frac{d\mathbb{Q}}{d\mathbb{P}}$, Radon-Nikodym derivative of $\mathbb{Q}$ with respect to $\mathbb{P}$.
- $U, V, W : \Omega \rightarrow \mathcal{P}(\mathbb{T})$, $\mathcal{H}/\mathcal{T}$ -measurable targeting functions.

Transition probability kernels

- J from $(\Omega, \mathcal{H})$ into $(\Omega, \mathcal{H}_{\mathbb{U}})$, an assignment kernel. Typically integrated out, and written $J(\eta; d\omega)$.
- $K_{\mathbb{S}}$ from $(\Omega, \mathcal{H}) \otimes (\Omega, \mathcal{H}_{\mathbb{S}})$ into $(\Omega, \mathcal{H})$ for $\mathbb{S} \in \mathcal{P}(\mathbb{T})$, causal kernels. Typically written $K_{\mathbb{S}}((\eta, \omega); A)$.
- $L_{\mathbb{S}}$ from $(\Omega, \mathcal{H}) \otimes (\Omega, \mathcal{H}_{\mathbb{S}})$ into $(\Omega, \mathcal{H}_{\mathbb{S} \cup \mathbb{U}})$, incoming causal kernels in soft interventions. Typically integrated out, and written $L_{\mathbb{S}}((\eta, \omega); d\omega')$.
- M, N , for other uses of transition probability kernels.

Homogenisations

- $\bar{J}$, homogenised assignment kernel.
- $\bar{K}_{\mathbb{U}}(\omega; A)$, homogenised uniform causal kernel.
- $\bar{K}_U((\eta, \omega); A)$, homogenised targeted causal kernel.
- $\bar{K}_{\mathbb{U}}^J(\omega; A)$, J -homogenisation of uniform causal kernels.
- $\bar{K}_U^J((\eta, \omega); A)$, J -homogenisations of targeted causal kernels.

Collections of transition probability kernels

- $\mathbb{K} = \{K_{\mathbb{S}} : \mathbb{S} \in \mathcal{P}(\mathbb{T})\}$, the causal mechanism.
- $\mathbb{L} = \{L_{\mathbb{S}} : \mathbb{S} \in \mathcal{P}(\mathbb{T})\}$, the incoming causal mechanism into $(\Omega, \mathcal{H}_{\mathbb{U}})$.

Interventional objects

- $\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}$, the interventional probability measure following a homogeneous intervention.
- $J^{\text{do}(\mathbb{U})}$, pointwise interventional probability measure.
- $\mathbb{P}^{\text{do}(\mathbb{U}, J)}$, the interventional probability measure following a uniform heterogeneous intervention.
- $\mathbb{P}^{\text{do}(U, J)}$, the interventional probability measure following a targeted intervention.

- $\mathbb{E}^{\text{do}(\mathbb{U}, \mathbb{Q})}$, $\mathbb{E}^{\text{do}(\mathbb{U}, J)}$ and $\mathbb{E}^{\text{do}(U, J)}$, expectations with respect to $\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}$, $\mathbb{P}^{\text{do}(\mathbb{U}, J)}$ and $\mathbb{P}^{\text{do}(U, J)}$.
- $K_{\mathbb{S}}^{\text{do}(\mathbb{U})}$, the interventional causal kernels following a hard intervention.
- $K_{\mathbb{S}}^{\text{do}(\mathbb{U}, \mathbb{L})}$, the interventional causal kernels following a soft intervention.
- $\mathbb{K}^{\text{do}(\mathbb{U})}$ and $\mathbb{K}^{\text{do}(\mathbb{U}, \mathbb{L})}$, the interventional causal mechanisms following hard and soft interventions.

Conditioning

- $\mathbb{P}_G$, the conditional probability given an event G .
- $\mathbb{P}_{\mathcal{G}}$, the conditional probability given a σ -algebra $\mathcal{G}$.
- $K_{\mathbb{S}|\mathcal{G}}((\eta, \omega)|\gamma; \cdot)$, the measure $K_{\mathbb{S}}((\eta, \omega); \cdot)$ conditioned on $\mathcal{G}$ at the outcome value γ .

Targeting-related objects

- $\mathcal{H}_U$, the targeted coordinate σ -algebra on $\Omega \times \Omega$.
- $\tilde{\mathcal{H}}_U = (\mathcal{H} \otimes \{\emptyset, \Omega\}) \vee \mathcal{H}_U$ on $\Omega \times \Omega$, the targeted coordinate σ -algebra $\mathcal{H}_U$ augmented with full information in the η -argument.
- $\tilde{\mathcal{H}}_U^{\mathcal{J}} = (\mathcal{J} \otimes \{\emptyset, \Omega\}) \vee \mathcal{H}_U$ on $\Omega \times \Omega$, the targeted coordinate σ -algebra $\mathcal{H}_U$ augmented with the $\mathcal{J}$ -information in the η -argument.
- $\Sigma_U(A)$ and Σ_U , heterogeneity σ -algebras of the targeted causal kernel K_U .

2 Preliminaries on probability theory

We collect some background on the theories of measures and probabilities needed to make this monograph self-contained. All proofs are provided for completeness. Readers encountering measure theory for the first time may find this chapter dense; for a comprehensive reference, we point to [Çinlar \(2011\)](#). Those who are familiar with measure-theoretic probability theory can skip to Chapter 3, and only refer to this chapter for specific results quoted in the later chapters.

We focus only on probability measures, and do not introduce measures in full generality. Hence, no consideration of concepts such as σ -finiteness will be needed. No exercises are provided for this chapter, unlike the subsequent chapters.

We begin with measurable spaces and measurable functions (Sections 2.1 and 2.2), then introduce probability measures (Section 2.3, including product probability measures in Section 2.4), integration/expectation (Section 2.5), the Radon-Nikodym theorem and Lebesgue decomposition (Section 2.6) and finally conditioning (Section 2.7). A strong emphasis will be placed on product structures and transition probability kernels, which form the backbone of causal space theory. In particular, we present in full the proofs of the Carathéodory extension theorem (Theorem 2.4.1) and the construction of product probability measures.

We denote by $\emptyset$ the empty set. For sets A, B , we denote their union, intersection, difference and symmetric difference by $A \cup B$, $A \cap B$, $A \setminus B$ and $A \triangle B$ respectively. Further, we write $A \sqcup B$ if we want to emphasise that the union is of disjoint sets. For any set A , we denote its power set, i.e. the set of all its subsets, by $\mathcal{P}(A) = 2^A$. Finally, we write $A \times B$ as the Cartesian product, i.e. the set of all possible ordered pairs (a, b) with $a \in A$ and $b \in B$.

We write $\mathbb{N} = \{1, 2, \dots\}$ for the set of natural numbers, $\mathbb{R} = (-\infty, \infty)$ for the set of real numbers, $\overline{\mathbb{R}} = [-\infty, \infty]$ for the set of real numbers extended to include $-\infty$ and ∞ , $\mathbb{R}_+ = \{a \in \mathbb{R} : a \geq 0\}$ for the set of non-negative real numbers and $\overline{\mathbb{R}}_+ = [0, \infty]$ for the set of non-negative real numbers extended to include ∞ . We use the convention that $0 \cdot \infty = 0$ and $\infty + a = \infty$ for $a \in \mathbb{R}$.

2.1 Measurable spaces

The starting point is a set Ω . For a subset $A \subseteq \Omega$, we denote its *complement* by $\Omega \setminus A$.

Suppose that $\mathcal{H}$ is a non-empty collection of subsets of Ω . We say that $\mathcal{H}$ is a σ -*algebra* if it is closed under complements and countable unions, i.e.

- (i) if $A \in \mathcal{H}$, then $\Omega \setminus A \in \mathcal{H}$; and
- (ii) if $A_1, A_2, \dots \in \mathcal{H}$, then $\cup_{n \in \mathbb{N}} A_n \in \mathcal{H}$.

We have the following two immediate consequences.

Lemma 2.1.1. *Suppose that $\mathcal{H}$ is a σ -algebra on Ω .*

- (i) $\emptyset, \Omega \in \mathcal{H}$.
- (ii) $\mathcal{H}$ is also closed under countable, hence finite, intersections.

Proof. (i) Indeed, since $\mathcal{H}$ is required to be non-empty, there is some $A \in \mathcal{H}$; then by closure under complements, $\Omega \setminus A \in \mathcal{H}$, and by closure under unions, $\Omega = A \cup (\Omega \setminus A) \in \mathcal{H}$. Then by closure under complements again, $\emptyset = \Omega \setminus \Omega \in \mathcal{H}$.

- (ii) Indeed, if $A_1, A_2, \dots \in \mathcal{H}$, then $\cap_{n \in \mathbb{N}} A_n = \Omega \setminus (\cup_{n \in \mathbb{N}} (\Omega \setminus A_n)) \in \mathcal{H}$, by closure under complements and countable unions. For a finite family $A_1, \dots, A_n \in \mathcal{H}$, set $A_m = \Omega \in \mathcal{H}$ for $m > n$, which lies in $\mathcal{H}$ by (i); then $\cap_{m=1}^n A_m = \cap_{m \in \mathbb{N}} A_m \in \mathcal{H}$. $\square$

Further, arbitrary non-empty (including uncountable) intersections of σ -algebras are σ -algebras.

Lemma 2.1.2. *Suppose that $(\mathcal{H}_t)_{t \in \mathbb{T}}$ is an arbitrary non-empty collection of σ -algebras on Ω . Then the intersection $\cap_{t \in \mathbb{T}} \mathcal{H}_t$ is also a σ -algebra on Ω .*

Proof. (i) First, note that $\Omega \in \cap_{t \in \mathbb{T}} \mathcal{H}_t$, so the intersection is not empty. Suppose that $A \in \cap_{t \in \mathbb{T}} \mathcal{H}_t$. Then $A \in \mathcal{H}_t$ for all $t \in \mathbb{T}$, meaning $\Omega \setminus A \in \mathcal{H}_t$ for all $t \in \mathbb{T}$ since each $\mathcal{H}_t$ is a σ -algebra. Hence, $\Omega \setminus A \in \cap_{t \in \mathbb{T}} \mathcal{H}_t$.

- (ii) Suppose that $A_1, A_2, \dots \in \cap_{t \in \mathbb{T}} \mathcal{H}_t$. Then $A_1, A_2, \dots \in \mathcal{H}_t$ for each $t \in \mathbb{T}$, meaning that $\cup_{n \in \mathbb{N}} A_n \in \mathcal{H}_t$ for all $t \in \mathbb{T}$ since $\mathcal{H}_t$ is a σ -algebra. Hence, $\cup_{n \in \mathbb{N}} A_n \in \cap_{t \in \mathbb{T}} \mathcal{H}_t$. $\square$

Starting from an arbitrary collection $\mathcal{C}$ of subsets of Ω , we define the σ -algebra *generated* by $\mathcal{C}$ to be the smallest σ -algebra containing $\mathcal{C}$. We denote it by $\sigma(\mathcal{C})$. Note that there is at least one such σ -algebra, namely, the power set $\mathcal{P}(\Omega)$, and then by Lemma 2.1.2, we can take the intersection of all σ -algebras containing $\mathcal{C}$.

Unlike intersections, unions of σ -algebras are in general not σ -algebras. For a collection $(\mathcal{H}_t)_{t \in \mathbb{T}}$ of σ -algebras on Ω , we denote by $\vee_{t \in \mathbb{T}} \mathcal{H}_t = \sigma(\cup_{t \in \mathbb{T}} \mathcal{H}_t)$ the σ -algebra generated by the union. The following lemma shows that it is in fact generated by the collection of *finite* intersections of sets drawn from the constituent σ -algebras.

Lemma 2.1.3. *Let $\mathbb{T}$ be a non-empty index set, and let $(\mathcal{H}_t)_{t \in \mathbb{T}}$ be a collection of σ -algebras on Ω . Define*

$$\mathcal{C} = \{\cap_{t \in \mathbb{S}} A_t : \mathbb{S} \subseteq \mathbb{T} \text{ finite and non-empty, and } A_t \in \mathcal{H}_t \text{ for each } t \in \mathbb{S}\}.$$

Then $\vee_{t \in \mathbb{T}} \mathcal{H}_t = \sigma(\mathcal{C})$.

Proof. For any finite non-empty $\mathbb{S} \subseteq \mathbb{T}$ and any $A_t \in \mathcal{H}_t$, $t \in \mathbb{S}$, we have $A_t \in \mathcal{H}_t \subseteq \bigvee_{t \in \mathbb{T}} \mathcal{H}_t$ for each $t \in \mathbb{S}$, and $\bigvee_{t \in \mathbb{T}} \mathcal{H}_t$ is a σ -algebra, hence closed under finite intersections by Lemma 2.1.1(ii). Therefore $\bigcap_{t \in \mathbb{S}} A_t \in \bigvee_{t \in \mathbb{T}} \mathcal{H}_t$, so that $\mathcal{C} \subseteq \bigvee_{t \in \mathbb{T}} \mathcal{H}_t$ and $\sigma(\mathcal{C}) \subseteq \bigvee_{t \in \mathbb{T}} \mathcal{H}_t$.

Conversely, for each $s \in \mathbb{T}$ and each $A \in \mathcal{H}_s$, taking $\mathbb{S} = \{s\}$ shows that $A \in \mathcal{C}$. Hence $\mathcal{H}_s \subseteq \mathcal{C}$ for every $s \in \mathbb{T}$, so that $\bigcup_{t \in \mathbb{T}} \mathcal{H}_t \subseteq \mathcal{C} \subseteq \sigma(\mathcal{C})$, and therefore $\bigvee_{t \in \mathbb{T}} \mathcal{H}_t = \sigma(\bigcup_{t \in \mathbb{T}} \mathcal{H}_t) \subseteq \sigma(\mathcal{C})$.

Combining the two inclusions gives the result. $\square$

An intersection $\bigcap_{t \in \mathbb{T}} A_t$ with $A_t \in \mathcal{H}_t$ over an uncountable index set need not belong to $\bigvee_{t \in \mathbb{T}} \mathcal{H}_t$.

For a σ -algebra $\mathcal{H}$ on Ω and an element $A \in \mathcal{H}$, we define the collection $\mathcal{H}|_A = \{A \cap B : B \in \mathcal{H}\}$. In the next result, we will show that in fact it is a σ -algebra on A . We call it the *trace* of $\mathcal{H}$ on A .

Lemma 2.1.4. $\mathcal{H}|_A$ is a σ -algebra on A .

Proof. (i) Suppose that $B \in \mathcal{H}|_A$. Then $B = A \cap C$ for some $C \in \mathcal{H}$, and since $\Omega \setminus C \in \mathcal{H}$, we have $A \cap (\Omega \setminus C) \in \mathcal{H}|_A$. But $A \cap (\Omega \setminus C) = A \setminus (A \cap C) = A \setminus B$, so $\mathcal{H}|_A$ is closed under complements.

(ii) Suppose that $B_1, B_2, \dots \in \mathcal{H}|_A$. Then there exist $C_1, C_2, \dots \in \mathcal{H}$ such that $B_n = C_n \cap A$ for each $n \in \mathbb{N}$, and $\bigcup_{n \in \mathbb{N}} C_n \in \mathcal{H}$. So

$$\bigcup_{n \in \mathbb{N}} B_n = \bigcup_{n \in \mathbb{N}} (A \cap C_n) = A \cap (\bigcup_{n \in \mathbb{N}} C_n) \in \mathcal{H}|_A,$$

hence $\mathcal{H}|_A$ is closed under countable unions. $\square$

A pair $(\Omega, \mathcal{H})$, where Ω is a set and $\mathcal{H}$ is a σ -algebra on Ω , is a *measurable space*. We call elements of a σ -algebra $\mathcal{H}$ *measurable sets*. The most important example of a measurable space is $\mathbb{R}$ (or $\overline{\mathbb{R}}$, $\mathbb{R}_+$ or $\overline{\mathbb{R}}_+$) equipped with the *Borel σ -algebra* $\mathcal{B}$, i.e. the σ -algebra generated by the open subsets.

There are other elementary collections of subsets in addition to σ -algebras. A collection of subsets of Ω is called a π -system if it is closed under finite intersections. A collection $\mathcal{D}$ of subsets of Ω is called a λ -system if the following are satisfied.

- $\Omega \in \mathcal{D}$.
- If $A, B \in \mathcal{D}$ and $A \supseteq B$, then $A \setminus B \in \mathcal{D}$.
- If $A_n \in \mathcal{D}$ for all $n \in \mathbb{N}$ and $A_1 \subseteq A_2 \subseteq \dots$, then $A = \bigcup_{n \in \mathbb{N}} A_n \in \mathcal{D}$.

In the last property, we call $(A_n)_{n \in \mathbb{N}}$ an *increasing sequence* with limit A . We now have the following useful result, to be invoked often throughout the monograph.

Theorem 2.1.5 (Monotone class theorem). *If a λ -system contains a π -system, then it also contains the σ -algebra generated by that π -system.*

Proof. Let $\mathcal{C}$ be a π -system. By a similar argument to the proof of Lemma 2.1.2, arbitrary intersections of λ -systems are λ -systems. So without loss of generality, we take $\mathcal{D}$ to be the λ -system that is the intersection of all λ -systems that contain $\mathcal{C}$. It suffices to show that $\mathcal{D}$ is a σ -algebra.

Firstly, $\mathcal{D}$ is closed under complements, since it is a λ -system. Closure under countable unions will be shown in several steps.

Step 1: $\mathcal{D}$ is closed under intersections. Take any $A \in \mathcal{D}$. Define $\mathcal{D}_A = \{B \in \mathcal{D} : A \cap B \in \mathcal{D}\}$. Then $\mathcal{D}_A$ is a λ -system, because

- clearly, $\Omega \in \mathcal{D}_A$, since $A \cap \Omega = A$;
- if $B, C \in \mathcal{D}_A$ and $B \supseteq C$, then $A \cap (B \setminus C) = (A \cap B) \setminus (A \cap C)$, where both $A \cap B$ and $A \cap C$ are in $\mathcal{D}$ since $B, C \in \mathcal{D}_A$, and since $\mathcal{D}$ is a λ -system, $A \cap (B \setminus C) \in \mathcal{D}$, which means $B \setminus C \in \mathcal{D}_A$; and
- if $B_n \in \mathcal{D}_A$ for all $n \in \mathbb{N}$ and $B_1 \subseteq B_2 \subseteq \dots$, then $A \cap (\cup_{n \in \mathbb{N}} B_n) = \cup_{n \in \mathbb{N}} (A \cap B_n)$, and since $B_n \in \mathcal{D}_A$, $A \cap B_n \in \mathcal{D}$ for all $n \in \mathbb{N}$, meaning $\cup_{n \in \mathbb{N}} (A \cap B_n) \in \mathcal{D}$, finally implying $\cup_{n \in \mathbb{N}} B_n \in \mathcal{D}_A$.

But if $A \in \mathcal{C}$, $\mathcal{D}_A$ clearly contains $\mathcal{C}$ since $\mathcal{C}$ is a π -system, so $\mathcal{D} \subseteq \mathcal{D}_A$, meaning that for all $A \in \mathcal{C}$ and all $B \in \mathcal{D}$, $A \cap B \in \mathcal{D}$. But this means that for all $B \in \mathcal{D}$, $\mathcal{D}_B$ contains $\mathcal{C}$, meaning $\mathcal{D} \subseteq \mathcal{D}_B$, which in turn implies that $\mathcal{D}$ is closed under intersections.

Step 2: $\mathcal{D}$ is closed under unions. Suppose $A, B \in \mathcal{D}$. Then $A \cup B = \Omega \setminus ((\Omega \setminus A) \cap (\Omega \setminus B))$, and since $\mathcal{D}$ is closed under complements (by the definition of a λ -system) and intersections (by Step 1), $A \cup B \in \mathcal{D}$.

Step 3: $\mathcal{D}$ is closed under countable unions. Suppose $A_1, A_2, \dots \in \mathcal{D}$. Then for each $n \in \mathbb{N}$, take $B_n = \cup_{m=1}^n A_m$. Then each B_n is in $\mathcal{D}$ by induction, since $B_1 \in \mathcal{D}$ and $B_n = B_{n-1} \cup A_n$, and we showed in Step 2 that $\mathcal{D}$ is closed under unions. But $(B_n)_{n \in \mathbb{N}}$ is an increasing sequence with limit $\cup_{n \in \mathbb{N}} A_n$, and since $\mathcal{D}$ is a λ -system, we must have $\cup_{n \in \mathbb{N}} A_n \in \mathcal{D}$. Hence $\mathcal{D}$ is closed under countable unions.

□

Yet another class of subset collections is called a *semi-ring*, which will be important in defining measures over product σ -algebras, via the Carathéodory extension theorem (see Theorem 2.4.1 in Section 2.4). A collection $\mathcal{S}$ of subsets of Ω is a semi-ring if the following are satisfied.

- $\emptyset \in \mathcal{S}$.
- If $A, B \in \mathcal{S}$, then $A \cap B \in \mathcal{S}$ (closure under intersection).
- If $A, B \in \mathcal{S}$, then there exist disjoint sets $S_i \in \mathcal{S}$, $i = 1, \dots, n$ for some $n \in \mathbb{N}$ such that

$$A \setminus B = \sqcup_{i=1}^n S_i.$$

In other words, the difference between any two sets can be decomposed as a finite disjoint union.

Further, a collection $\mathcal{R}$ of subsets of Ω is an *algebra* if $\Omega \in \mathcal{R}$, and it is closed under complements, finite unions and finite intersections.

In causal space theory, product structures of measurable spaces will play a crucial role, with each component of the product playing the role of a “variable” or a “coordinate”. Let $\mathbb{T}$ be an arbitrary index set, allowed to be finite, countably infinite or uncountably infinite. For each $t \in \mathbb{T}$, suppose that we have a measurable space $(\Omega_t, \mathcal{E}_t)$. A *measurable cylinder* in $\times_{t \in \mathbb{T}} \Omega_t$ is a subset

$$\times_{t \in \mathbb{T}} A_t = \{(\omega_t)_{t \in \mathbb{T}} : \omega_t \in A_t\},$$

where $A_t \in \mathcal{E}_t$ for each $t \in \mathbb{T}$ and differs from Ω_t for only a finite number of $t \in \mathbb{T}$. The collection of all measurable cylinders is a prototypical example of a semi-ring:

Lemma 2.1.6. *The collection $\mathcal{S}$ of all measurable cylinders is a semi-ring.*

Proof. Clearly, $\emptyset, \Omega \in \mathcal{S}$. If $C = \times_{t \in \mathbb{T}} A_t$ and $D = \times_{t \in \mathbb{T}} B_t$ are measurable cylinders, then $C \cap D = \times_{t \in \mathbb{T}} (A_t \cap B_t)$, which is again a measurable cylinder. It remains to check finite decomposability of differences. Let $\mathbb{S} = \{t_1, \dots, t_n\} \subseteq \mathbb{T}$ be a finite set containing all coordinates on which either C or D differs from the full coordinate space. Then

$$C \setminus D = \sqcup_{j=1}^n [(\times_{i < j} (A_{t_i} \cap B_{t_i})) \times (A_{t_j} \setminus B_{t_j}) \times (\times_{i > j} A_{t_i}) \times (\times_{t \notin \mathbb{S}} \Omega_t)],$$

after deleting any empty terms. Each term is a measurable cylinder, and the union is disjoint. Hence $\mathcal{S}$ is a semi-ring. $\square$

The *product σ -algebra* $\mathcal{H} = \otimes_{t \in \mathbb{T}} \mathcal{E}_t$ on $\Omega = \times_{t \in \mathbb{T}} \Omega_t$ is the σ -algebra generated by the collection of all measurable cylinders. Then the *product measurable space* is $(\Omega, \mathcal{H}) = \otimes_{t \in \mathbb{T}} (\Omega_t, \mathcal{E}_t)$.

For each subset $\mathbb{S} \in \mathcal{P}(\mathbb{T})$, we write $\Omega_{\mathbb{S}} = \times_{s \in \mathbb{S}} \Omega_s$, and for each $\omega = (\omega_t)_{t \in \mathbb{T}}$, we write $\omega_{\mathbb{S}} = (\omega_s)_{s \in \mathbb{S}} \in \Omega_{\mathbb{S}}$, where $\omega_s \in \Omega_s$ for each $s \in \mathbb{S}$. We define $\mathcal{H}_{\mathbb{S}}$ as the sub- σ -algebra of $\mathcal{H}$ on Ω generated by measurable cylinders $\times_{t \in \mathbb{T}} A_t$, where $A_t \in \mathcal{E}_t$ for all $t \in \mathbb{T}$, $A_t = \Omega_t$ for all $t \notin \mathbb{S}$ and $A_t \neq \Omega_t$ for only finitely many $t \in \mathbb{S}$. In particular, $\mathcal{H}_{\emptyset} = \{\emptyset, \Omega\}$ and $\mathcal{H}_{\mathbb{T}} = \mathcal{H}$. Note also that, if $\mathbb{S} \subseteq \mathbb{U}$, then $\mathcal{H}_{\mathbb{S}} \subseteq \mathcal{H}_{\mathbb{U}}$, and that $\mathcal{H}_{\mathbb{S}} \vee \mathcal{H}_{\mathbb{U}} = \mathcal{H}_{\mathbb{S} \cup \mathbb{U}}$.

We also define $\mathcal{E}_{\mathbb{S}} = \otimes_{t \in \mathbb{S}} \mathcal{E}_t$ as the σ -algebra on $\Omega_{\mathbb{S}}$ generated by measurable cylinders $\times_{t \in \mathbb{S}} A_t$. Note that $\mathcal{H}_{\mathbb{S}}$ and $\mathcal{E}_{\mathbb{S}}$ are different objects: the former is a σ -algebra on $\Omega = \times_{t \in \mathbb{T}} \Omega_t$, and is a sub- σ -algebra of $\mathcal{H}$, whereas the latter is a σ -algebra on $\Omega_{\mathbb{S}}$, and is *not* a sub- σ -algebra of $\mathcal{H}$.

2.2 Measurable functions

Let $(\Omega_1, \mathcal{H}_1)$ and $(\Omega_2, \mathcal{H}_2)$ be measurable spaces. We say that a function $f : \Omega_1 \rightarrow \Omega_2$ is $\mathcal{H}_1/\mathcal{H}_2$ -*measurable* if, for all $B \in \mathcal{H}_2$, the pre-image

$$f^{-1}(B) = \{\omega \in \Omega_1 : f(\omega) \in B\}$$

is in $\mathcal{H}_1$. For some collection $\mathcal{C}$ of subsets of Ω_2 , let us denote $f^{-1}(\mathcal{C}) = \{f^{-1}(B) : B \in \mathcal{C}\}$. Then, $f^{-1}(\sigma(\mathcal{C}))$ is a σ -algebra, since, for any $A \in f^{-1}(\sigma(\mathcal{C}))$ with $A = f^{-1}(B)$, $\Omega \setminus A = f^{-1}(\Omega \setminus B) \in f^{-1}(\sigma(\mathcal{C}))$, and for $A_n \in f^{-1}(\sigma(\mathcal{C}))$ with $A_n = f^{-1}(B_n)$, $\cup_{n \in \mathbb{N}} A_n = f^{-1}(\cup_{n \in \mathbb{N}} B_n) \in f^{-1}(\sigma(\mathcal{C}))$. Hence, we also have $\sigma(f^{-1}(\mathcal{C})) = f^{-1}(\sigma(\mathcal{C}))$.

The following lemma characterises measurability using generators.

Lemma 2.2.1. *If $\mathcal{H}_2$ is generated by some collection $\mathcal{C}$, then f is $\mathcal{H}_1/\mathcal{H}_2$ -measurable if and only if $f^{-1}(\mathcal{C}) \subseteq \mathcal{H}_1$.*

Proof. The “only if” part is obvious. Now suppose that $f^{-1}(B) \in \mathcal{H}_1$ for all $B \in \mathcal{C}$. Define a collection $\mathcal{M} = \{A \in \mathcal{H}_2 : f^{-1}(A) \in \mathcal{H}_1\}$. Then clearly, $\mathcal{M} \subseteq \mathcal{H}_2$ and $\mathcal{C} \subseteq \mathcal{M}$. We will show that $\mathcal{M}$ is a σ -algebra, which would imply that $\mathcal{H}_2 \subseteq \mathcal{M}$, since $\mathcal{H}_2$ is generated by $\mathcal{C}$, and hence the smallest σ -algebra containing it. Combining the two inclusions would imply $\mathcal{M} = \mathcal{H}_2$, and so $f^{-1}(A) \in \mathcal{H}_1$ for all $A \in \mathcal{H}_2$, meaning f would be $\mathcal{H}_1/\mathcal{H}_2$ -measurable.

First note that $\mathcal{M}$ is non-empty, since $f^{-1}(\Omega_2) = \Omega_1 \in \mathcal{H}_1$ so $\Omega_2 \in \mathcal{M}$. Also, if $A \in \mathcal{M}$, then $f^{-1}(\Omega_2 \setminus A) = \Omega_1 \setminus f^{-1}(A) \in \mathcal{H}_1$, since $f^{-1}(A) \in \mathcal{H}_1$ and $\mathcal{H}_1$ is a σ -algebra. Finally, if $A_1, A_2, \dots \in \mathcal{M}$, then $f^{-1}(\cup_{n \in \mathbb{N}} A_n) = \cup_{n \in \mathbb{N}} f^{-1}(A_n) \in \mathcal{H}_1$, since $f^{-1}(A_n) \in \mathcal{H}_1$ for each $n \in \mathbb{N}$ and $\mathcal{H}_1$ is a σ -algebra. Hence $\mathcal{M}$ is a σ -algebra, and the proof is complete. $\square$

For a set Ω_1 , a measurable space $(\Omega_2, \mathcal{H}_2)$ and a function $f : \Omega_1 \rightarrow \Omega_2$, the σ -algebra generated by f is the smallest σ -algebra on Ω_1 that makes f measurable, and is denoted as $\sigma(f)$. Note that there is at least one such σ -algebra, namely the power set $\mathcal{P}(\Omega_1)$, and we can take the intersection of all such σ -algebras by Lemma 2.1.2.

If $(\Omega_1, \mathcal{H}_1)$, $(\Omega_2, \mathcal{H}_2)$ and $(\Omega_3, \mathcal{H}_3)$ are measurable spaces, and $f : \Omega_1 \rightarrow \Omega_2$ and $g : \Omega_2 \rightarrow \Omega_3$ are measurable functions, then it is clear that the composition $g \circ f : \Omega_1 \rightarrow \Omega_3$ is $\mathcal{H}_1/\mathcal{H}_3$ -measurable. Indeed, for any $A \in \mathcal{H}_3$, $g^{-1}(A) \in \mathcal{H}_2$ by the $\mathcal{H}_2/\mathcal{H}_3$ -measurability of g , and hence $(g \circ f)^{-1}(A) = f^{-1}(g^{-1}(A)) \in \mathcal{H}_1$ by the $\mathcal{H}_1/\mathcal{H}_2$ -measurability of f .

The following lemma collects some elementary properties of measurable functions.

Lemma 2.2.2. *Let $f, g : \Omega \rightarrow \overline{\mathbb{R}}$ be $\mathcal{H}$ -measurable, and $b \in \mathbb{R}$. Then bf , $f + g$, $f - g$, $\max\{f, g\}$, and $\min\{f, g\}$ are $\mathcal{H}$ -measurable, provided that they are well-defined.*

Before we begin with the proof, we present some shorthand notation. First, note that, for any $a \in \mathbb{R}$, the intervals $(a, \infty]$, $[a, \infty]$, $[-\infty, a)$ and $[-\infty, a]$ clearly belong to $\mathcal{B}$. Then for a function $f : \Omega \rightarrow \overline{\mathbb{R}}$, we write $\{f > a\} = f^{-1}((a, \infty])$, and similarly for $\{f < a\}$, $\{f \leq a\}$ and $\{f \geq a\}$. Thus, if f is $\mathcal{H}$ -measurable, then we have $\{f > a\} \in \mathcal{H}$, and likewise for $\{f < a\}$, $\{f \leq a\}$ and $\{f \geq a\}$. Conversely, since the Borel σ -algebra $\mathcal{B}$ is generated by the intervals $(a, \infty]$ for $a \in \mathbb{R}$, if $\{f > a\} \in \mathcal{H}$ for all $a \in \mathbb{R}$, then by Lemma 2.2.1, f is $\mathcal{H}$ -measurable, and likewise for $\{f < a\}$, $\{f \leq a\}$ and $\{f \geq a\}$.

Proof. Scalar product. For any $a \in \mathbb{R}$, $\{bf < a\} = \{f < \frac{a}{b}\}$ if $b \neq 0$. If $b = 0$, then the constant function bf is clearly measurable.

Sum. For any $a \in \mathbb{R}$,

$$\{f + g < a\} = \bigcup_{q \text{ rational}} (\{f < q\} \cap \{g < a - q\}).$$

Indeed, if $f(\omega) < q$ and $g(\omega) < a - q$, then $f(\omega) + g(\omega) < a$; conversely, if $f(\omega) + g(\omega) < a$, then by density of rational numbers in $\mathbb{R}$ there exists a rational number q with $f(\omega) < q < a - g(\omega)$, so that $f(\omega) < q$ and $g(\omega) < a - q$. The right-hand side is a countable union of sets in $\mathcal{H}$, and so $\{f + g < a\} \in \mathcal{H}$.

Difference. First, $-g$ is $\mathcal{H}$ -measurable, since for any $a \in \mathbb{R}$, $\{-g < a\} = \{g > -a\} = \Omega \setminus \{g \leq -a\} \in \mathcal{H}$ by measurability of g and closure of $\mathcal{H}$ under complements. Hence $f - g = f + (-g)$ is $\mathcal{H}$ -measurable by the previous part.

Maximum and minimum. For any $a \in \mathbb{R}$,

$$\{\max\{f, g\} < a\} = \{f < a\} \cap \{g < a\} \in \mathcal{H},$$

and

$$\{\min\{f, g\} < a\} = \{f < a\} \cup \{g < a\} \in \mathcal{H},$$

as required. $\square$

Later in Lemma 2.2.7, the product fg will also be shown to be measurable. For any $\mathcal{H}$ -measurable function $f : \Omega \rightarrow \overline{\mathbb{R}}$, let us define the positive and negative parts of f as $f_+ = \max\{f, 0\}$ and $f_- = -\min\{f, 0\}$, which are both measurable (by Lemma 2.2.2) and take values in $\mathbb{R}_+$.

Suppose that $(\Omega, \mathcal{H}) = \otimes_{t \in \mathbb{T}} (\Omega_t, \mathcal{E}_t)$ is a product measurable space, as in Section 2.1. For any $\mathbb{S} \subseteq \mathbb{T}$, we define the *canonical projection* $\pi_{\mathbb{S}} : \Omega \rightarrow \Omega_{\mathbb{S}}$ by $\pi_{\mathbb{S}}(\omega) = \omega_{\mathbb{S}}$, which is $\mathcal{H}/\mathcal{E}_{\mathbb{S}}$ -measurable since, for any $A \in \mathcal{E}_{\mathbb{S}}$, the pre-image $\pi_{\mathbb{S}}^{-1}(A) = A \times \Omega_{\mathbb{T} \setminus \mathbb{S}}$ is a measurable cylinder in the generating semi-ring of $\mathcal{H} = \mathcal{E}_{\mathbb{S}} \otimes \mathcal{E}_{\mathbb{T} \setminus \mathbb{S}}$. Then the σ -algebras $\mathcal{H}_{\mathbb{S}}$ and $\mathcal{E}_{\mathbb{S}}$ are related by $\mathcal{H}_{\mathbb{S}} = \pi_{\mathbb{S}}^{-1}(\mathcal{E}_{\mathbb{S}})$, and $\mathcal{H}_{\mathbb{S}} = \sigma(\pi_{\mathbb{S}})$.

For a fixed $\omega_{\mathbb{S}} \in \Omega_{\mathbb{S}}$, let us also define the *insertion map* $\iota_{\omega_{\mathbb{S}}} : \Omega_{\mathbb{T} \setminus \mathbb{S}} \rightarrow \Omega$ by $\iota_{\omega_{\mathbb{S}}}(\omega_{\mathbb{T} \setminus \mathbb{S}}) = (\omega_{\mathbb{S}}, \omega_{\mathbb{T} \setminus \mathbb{S}})$. Then this is $\mathcal{E}_{\mathbb{T} \setminus \mathbb{S}}/\mathcal{H}$ -measurable. Indeed, for any measurable cylinder $\times_{t \in \mathbb{T}} A_t$ in the generating semi-ring of $\mathcal{H} = \otimes_{t \in \mathbb{T}} \mathcal{E}_t$, we have $\iota_{\omega_{\mathbb{S}}}^{-1}(\times_{t \in \mathbb{T}} A_t) = \emptyset$ if $\omega_s \notin A_s$ for some $s \in \mathbb{S}$, and $\iota_{\omega_{\mathbb{S}}}^{-1}(\times_{t \in \mathbb{T}} A_t) = \times_{t \in \mathbb{T} \setminus \mathbb{S}} A_t$ if $\omega_s \in A_s$ for all $s \in \mathbb{S}$. In both cases, we have $\iota_{\omega_{\mathbb{S}}}^{-1}(\times_{t \in \mathbb{T}} A_t) \in \mathcal{E}_{\mathbb{T} \setminus \mathbb{S}}$. Hence, by Lemma 2.2.1, the insertion map $\iota_{\omega_{\mathbb{S}}}$ is $\mathcal{E}_{\mathbb{T} \setminus \mathbb{S}}/\mathcal{H}$ -measurable.

For a fixed $A \in \mathcal{H}$ and an element $\omega_{\mathbb{S}} \in \Omega_{\mathbb{S}}$, we define the *section* of A at $\omega_{\mathbb{S}}$ as

$$A_{\omega_{\mathbb{S}}} = \{\omega_{\mathbb{T} \setminus \mathbb{S}} \in \Omega_{\mathbb{T} \setminus \mathbb{S}} : (\omega_{\mathbb{S}}, \omega_{\mathbb{T} \setminus \mathbb{S}}) \in A\} \subseteq \Omega_{\mathbb{T} \setminus \mathbb{S}}. \quad (2.1)$$

The section $A_{\omega_{\mathbb{S}}}$ is precisely the pre-image $\iota_{\omega_{\mathbb{S}}}^{-1}(A)$, and hence belongs to $\mathcal{E}_{\mathbb{T} \setminus \mathbb{S}}$.

Further to the product space $(\Omega, \mathcal{H}) = \otimes_{t \in \mathbb{T}} (\Omega_t, \mathcal{E}_t)$, let $(\Omega', \mathcal{H}')$ be another measurable space. Then we have the following lemma about what it means for a function to be measurable with respect to the coordinate σ -algebra $\mathcal{H}_{\mathbb{S}}$ for some $\mathbb{S} \subseteq \mathbb{T}$.

Lemma 2.2.3. *Suppose that $(\Omega', \mathcal{H}')$ is a measurable space such that $\mathcal{H}'$ separates points, in the sense that for any $\omega'_1, \omega'_2 \in \Omega'$ with $\omega'_1 \neq \omega'_2$, there exists $B \in \mathcal{H}'$ with $\omega'_1 \in B$ and $\omega'_2 \notin B$. Then for an $\mathcal{H}/\mathcal{H}'$ -measurable function $f : \Omega \rightarrow \Omega'$, the following are equivalent.*

(i) f is $\mathcal{H}_{\mathbb{S}}/\mathcal{H}'$ -measurable.

(ii) f depends only on the $\mathbb{S}$ -coordinates, in the sense that $f(\omega_1) = f(\omega_2)$ for all $\omega_1, \omega_2 \in \Omega$ with $\pi_{\mathbb{S}}(\omega_1) = \pi_{\mathbb{S}}(\omega_2)$.

(iii) There exists an $\mathcal{E}_{\mathbb{S}}/\mathcal{H}'$ -measurable function $h : \Omega_{\mathbb{S}} \rightarrow \Omega'$ such that $f = h \circ \pi_{\mathbb{S}}$.

Proof. **(i) $\implies$ (ii)** Let $\omega_1, \omega_2 \in \Omega$ with $\pi_{\mathbb{S}}(\omega_1) = \pi_{\mathbb{S}}(\omega_2)$. If $f(\omega_1) \neq f(\omega_2)$, then since $\mathcal{H}'$ separates points, there is $B \in \mathcal{H}'$ with $f(\omega_1) \in B$ and $f(\omega_2) \notin B$. By (i) and $\mathcal{H}_{\mathbb{S}} = \pi_{\mathbb{S}}^{-1}(\mathcal{E}_{\mathbb{S}})$, we have $f^{-1}(B) = \pi_{\mathbb{S}}^{-1}(E)$ for some $E \in \mathcal{E}_{\mathbb{S}}$. Then $\omega_1 \in f^{-1}(B)$ gives $\pi_{\mathbb{S}}(\omega_1) \in E$, hence $\pi_{\mathbb{S}}(\omega_2) \in E$ and so $\omega_2 \in f^{-1}(B)$. This is a contradiction.

(ii) $\implies$ (iii) Set $h = f \circ \iota_{\omega_{\mathbb{T} \setminus \mathbb{S}}} : \Omega_{\mathbb{S}} \rightarrow \Omega'$ for any $\omega_{\mathbb{T} \setminus \mathbb{S}} \in \Omega_{\mathbb{T} \setminus \mathbb{S}}$, which is $\mathcal{E}_{\mathbb{S}}/\mathcal{H}'$ -measurable as a composition of measurable functions. Then for any $\omega \in \Omega$, the outcome $\iota_{\omega_{\mathbb{T} \setminus \mathbb{S}}}(\pi_{\mathbb{S}}(\omega))$ has the same $\mathbb{S}$ -coordinates as ω , so (ii) gives $h(\pi_{\mathbb{S}}(\omega)) = f(\iota_{\omega_{\mathbb{T} \setminus \mathbb{S}}}(\pi_{\mathbb{S}}(\omega))) = f(\omega)$. Hence $f = h \circ \pi_{\mathbb{S}}$.

(iii) $\implies$ (i) Since $\mathcal{H}_{\mathbb{S}} = \sigma(\pi_{\mathbb{S}})$, the projection $\pi_{\mathbb{S}}$ is $\mathcal{H}_{\mathbb{S}}/\mathcal{E}_{\mathbb{S}}$ -measurable, so $f = h \circ \pi_{\mathbb{S}}$ is $\mathcal{H}_{\mathbb{S}}/\mathcal{H}'$ -measurable as a composition of measurable functions. $\square$

Further, for a function that is measurable with respect to the product σ -algebra, holding some of the coordinates fixed yields a function that is measurable with respect to the σ -algebra corresponding to the remaining coordinates. We continue to work with a product space $(\Omega, \mathcal{H}) = \otimes_{t \in \mathbb{T}} (\Omega_t, \mathcal{E}_t)$ and another measurable space $(\Omega', \mathcal{H}')$.

Lemma 2.2.4. *For any $\mathcal{H}/\mathcal{H}'$ -measurable function $f : \Omega \rightarrow \Omega'$, any $\mathbb{S} \in \mathcal{P}(\mathbb{T})$ and any $\omega_{\mathbb{T} \setminus \mathbb{S}} \in \Omega_{\mathbb{T} \setminus \mathbb{S}}$, the map $f_{\omega_{\mathbb{T} \setminus \mathbb{S}}} : \Omega_{\mathbb{S}} \rightarrow \Omega'$ defined by $f_{\omega_{\mathbb{T} \setminus \mathbb{S}}}(\omega_{\mathbb{S}}) = f((\omega_{\mathbb{S}}, \omega_{\mathbb{T} \setminus \mathbb{S}}))$ is $\mathcal{E}_{\mathbb{S}}$ -measurable.*

Proof. Take any $B \in \mathcal{H}'$. Then $f_{\omega_{\mathbb{T} \setminus \mathbb{S}}}^{-1}(B)$ is the section of $f^{-1}(B)$ at $\omega_{\mathbb{T} \setminus \mathbb{S}}$ (Eq. (2.1)), so is $\mathcal{E}_{\mathbb{S}}$ -measurable. $\square$

If Ω_2 is a subset of $\overline{\mathbb{R}}$ and $\mathcal{H}_2$ is the trace Borel σ -algebra, then we simply say that f is $\mathcal{H}_1$ -measurable if it is $\mathcal{H}_1/\mathcal{B}|_{\Omega_2}$ -measurable. If the σ -algebras under consideration are clear from context, then we simply say that a function is measurable.

Two elementary types of measurable functions are of particular interest. Suppose that we have a measurable space $(\Omega, \mathcal{H})$. Then for each $A \in \mathcal{H}$, we define the *indicator function* $\mathbf{1}_A : \Omega \rightarrow \{0, 1\}$ as

$$\mathbf{1}_A(\omega) = \begin{cases} 1 & \text{if } \omega \in A \\ 0 & \text{otherwise.} \end{cases} \quad (2.2)$$

The indicator function $\mathbf{1}_A$ is $\mathcal{H}$ -measurable, because $\mathbf{1}_A^{-1}(B)$ is one of $\emptyset, A, \Omega \setminus A, \Omega$, depending on which of 0 and 1 belong to B .

A function $f : \Omega \rightarrow \mathbb{R}$ is called a *simple* function if it has the form

$$f = \sum_{i=1}^n a_i \mathbf{1}_{A_i} \quad (2.3)$$

for some $n \in \mathbb{N}$, $a_i \in \mathbb{R}$ and $A_i \in \mathcal{H}$ for $i = 1, \dots, n$. We say that a simple function f is in *canonical form* if the $(A_i)_{i=1}^n$ form a disjoint partition of Ω and all a_i are distinct. We show that a canonical form always exists.

Lemma 2.2.5. *If $f = \sum_{i=1}^n a_i \mathbf{1}_{A_i}$ is a simple function, then there exist a disjoint partition $(B_j)_{j=1}^m$ of Ω and distinct real numbers $b_j \in \mathbb{R}$ such that $f = \sum_{j=1}^m b_j \mathbf{1}_{B_j}$.*

Proof. For each $\epsilon = (\epsilon_1, \dots, \epsilon_n) \in \{0, 1\}^n$, define

$$A'_\epsilon = (\cap_{i:\epsilon_i=1} A_i) \cap (\cap_{i:\epsilon_i=0} (\Omega \setminus A_i)).$$

Then $A'_\epsilon \in \mathcal{H}$ for each $\epsilon \in \{0, 1\}^n$ and $\{A'_\epsilon : \epsilon \in \{0, 1\}^n\}$ is a disjoint partition of Ω . For $\omega \in A'_\epsilon$, we have $f(\omega) = \sum_{i:\epsilon_i=1} a_i$. Hence f takes at most 2^n distinct values. Let $b_1, \dots, b_m$ be those values without repetition, and for each $j = 1, \dots, m$, define $B_j = \{\omega \in \Omega : f(\omega) = b_j\}$. Then each B_j is a union of some A'_ϵ , hence $B_j \in \mathcal{H}$ for all $j = 1, \dots, m$, and $f = \sum_{j=1}^m b_j \mathbf{1}_{B_j}$ is a canonical form of f . $\square$

Simple functions are also $\mathcal{H}$ -measurable, since, for every Borel set $A \subseteq \mathbb{R}$, $f^{-1}(A)$ is a finite union of some of the sets B_j , where B_j are from the canonical representation of f . If f and g are simple functions, it is immediate that $f + g$, $f - g$ and fg are simple functions, as well as af for any $a \in \mathbb{R}$.

A simple function f is said to be *positive* if $f(\omega) \geq 0$ for all $\omega \in \Omega$, equivalently if, in its canonical form (Lemma 2.2.5), we have $a_i \geq 0$ for all $i = 1, \dots, n$. An important result is the fact that we can approximate any non-negative measurable function as a limit of an increasing sequence of positive simple functions.

Theorem 2.2.6. (i) *Let $(f_n)_{n \in \mathbb{N}}$, $f_n : \Omega \rightarrow \overline{\mathbb{R}}_+$, be a sequence of $\mathcal{H}$ -measurable functions. If $\lim_{n \rightarrow \infty} f_n$ exists (which may be infinite), this limit is also $\mathcal{H}$ -measurable.*

(ii) *A function $f : \Omega \rightarrow \overline{\mathbb{R}}_+$ is $\mathcal{H}$ -measurable if and only if it is the pointwise limit of an increasing sequence of positive simple functions.*

Proof. (i) Write $f = \lim_{n \rightarrow \infty} f_n$. We first show that countable suprema and infima of measurable functions are measurable. For every $a \in \mathbb{R}$, $\{\sup_{n \in \mathbb{N}} f_n > a\} = \cup_{n \in \mathbb{N}} \{f_n > a\}$, which belongs to $\mathcal{H}$ by closure under countable unions. Hence $\sup_{n \in \mathbb{N}} f_n$ is $\mathcal{H}$ -measurable. Similarly, for every $a \in \mathbb{R}$, $\{\inf_{n \in \mathbb{N}} f_n < a\} = \cup_{n \in \mathbb{N}} \{f_n < a\}$, which belongs to $\mathcal{H}$, and so $\inf_{n \in \mathbb{N}} f_n$ is $\mathcal{H}$ -measurable.

Hence, $\limsup_{n \rightarrow \infty} f_n = \inf_{m \in \mathbb{N}} \sup_{n \geq m} f_n$ and $\liminf_{n \rightarrow \infty} f_n = \sup_{m \in \mathbb{N}} \inf_{n \geq m} f_n$ are $\mathcal{H}$ -measurable. Since $f = \lim_{n \rightarrow \infty} f_n$ exists, we have

$$f = \limsup_{n \rightarrow \infty} f_n = \liminf_{n \rightarrow \infty} f_n,$$

and so f is $\mathcal{H}$ -measurable.

(ii) Sufficiency follows from part (i). For necessity, for each $n \in \mathbb{N}$, define

$$d_n(a) = \sum_{k=1}^{n2^n} \frac{k-1}{2^n} \mathbf{1}_{[\frac{k-1}{2^n}, \frac{k}{2^n})}(a) + n \mathbf{1}_{[n, \infty)}(a), \quad a \in \overline{\mathbb{R}}_+, \quad (2.4)$$

and let $f_n = d_n \circ f$. In words, d_n maps any $a \geq n$ to n . Further, it divides the interval $[0, n)$ into equal intervals of length $\frac{1}{2^n}$ (so the intervals become finer with increasing n), and maps any $a < n$ to the lower end of the interval it belongs to. Clearly, each d_n is a positive simple function (in its canonical form), and hence so is each f_n . Further, we have $\lim_{n \rightarrow \infty} f_n(\omega) = f(\omega)$ for each $\omega \in \Omega$, since, for all $n \geq f(\omega)$, $f_n(\omega) = \frac{\lfloor 2^n f(\omega) \rfloor}{2^n}$, which converges to $f(\omega)$ as $n \rightarrow \infty$, and if $f(\omega) = \infty$, then $f_n(\omega) = n \rightarrow \infty = f(\omega)$. Finally, $(f_n)_{n \in \mathbb{N}}$ is clearly an increasing sequence, since for each $\omega \in \Omega$,

$$f_{n-1}(\omega) = \frac{\lfloor 2^{n-1} f(\omega) \rfloor}{2^{n-1}} \leq \frac{\lfloor 2^n f(\omega) \rfloor}{2^n} = f_n(\omega)$$

if $n-1 > f(\omega)$, and $f_n(\omega) \geq n-1 = f_{n-1}(\omega)$ otherwise. $\square$

Now, as promised, we show that the product of two measurable functions is measurable.

Lemma 2.2.7. *Let $f, g : \Omega \rightarrow \overline{\mathbb{R}}$ be $\mathcal{H}$ -measurable. Then fg is $\mathcal{H}$ -measurable.*

Proof. First suppose that $f, g : \Omega \rightarrow \overline{\mathbb{R}}_+$. For each $m \in \mathbb{N}$, the functions $d_m \circ f$ and $d_m \circ g$ are positive simple functions, so their product $(d_m \circ f)(d_m \circ g)$ is a positive simple function, and hence $\mathcal{H}$ -measurable. The sequence $((d_m \circ f)(d_m \circ g))_{m \in \mathbb{N}}$ is increasing, since $(d_m \circ f)_{m \in \mathbb{N}}$ and $(d_m \circ g)_{m \in \mathbb{N}}$ are increasing and non-negative. It converges pointwise to fg : this is clear where f and g are both finite; if $f(\omega) = \infty$ and $g(\omega) = 0$, both sides are 0 by the convention $0 \cdot \infty = 0$; if $f(\omega) = \infty$ and $g(\omega) > 0$, then $d_m \circ f(\omega) = m$ while $d_m \circ g(\omega) \rightarrow g(\omega) > 0$, so the product tends to ∞ ; and the remaining cases follow by interchanging the roles of f and g . By Theorem 2.2.6(i), fg is $\mathcal{H}$ -measurable.

For general $f, g : \Omega \rightarrow \overline{\mathbb{R}}$, write $fg = f_+g_+ - f_+g_- - f_-g_+ + f_-g_-$, which is well-defined since at each $\omega \in \Omega$, at most one of the four terms is non-zero, and where each of the four terms is $\mathcal{H}$ -measurable by the above. Now apply Lemma 2.2.2 repeatedly. $\square$

A particularly nice class of measurable spaces is called the *standard measurable spaces*, often also referred to as standard Borel measurable spaces. Suppose that $(\Omega_1, \mathcal{H}_1)$ and $(\Omega_2, \mathcal{H}_2)$ are measurable spaces. Then we say that a function $f : \Omega_1 \rightarrow \Omega_2$ is an *isomorphism* if it is a bijection, and both f and f^{-1} are $\mathcal{H}_1/\mathcal{H}_2$ -measurable. We say that a measurable space $(\Omega, \mathcal{H})$ is *standard* if it is isomorphic to a subset of $\mathbb{R}$ that belongs to the Borel σ -algebra $\mathcal{B}$, with its trace σ -algebra.

2.3 Probability measures and probability spaces

In this section, we shift our focus to probability theory. Henceforth, in the measurable space $(\Omega, \mathcal{H})$, we will refer to elements $\omega \in \Omega$ as *outcomes*, and sets $A \in \mathcal{H}$ as *events*.

A function $\mathbb{P} : \mathcal{H} \rightarrow [0, 1]$ is called a *probability measure* if

- (i) $\mathbb{P}(\Omega) = 1$; and
- (ii) $\mathbb{P}(\sqcup_{n \in \mathbb{N}} A_n) = \sum_{n \in \mathbb{N}} \mathbb{P}(A_n)$ for all disjoint sequences $(A_n)_{n \in \mathbb{N}}$ in $\mathcal{H}$, i.e. $\mathbb{P}$ is *σ -additive*.

A triple $(\Omega, \mathcal{H}, \mathbb{P})$, where $(\Omega, \mathcal{H})$ is a measurable space and $\mathbb{P}$ is a probability measure, is called a *probability space*. For each event $A \in \mathcal{H}$, we interpret $\mathbb{P}(A)$ as the *probability* that the event A occurs. Given a probability space $(\Omega, \mathcal{H}, \mathbb{P})$ and another measurable space $(\Omega_2, \mathcal{H}_2)$, an $\mathcal{H}/\mathcal{H}_2$ -measurable function $f : \Omega \rightarrow \Omega_2$ is called a *random variable*.

Throughout the monograph, we will often say that a property holds *$\mathbb{P}$ -almost surely*, or more precisely, for *$\mathbb{P}$ -almost all* $\omega \in \Omega$. This means that the set of all $\omega \in \Omega$ for which the property does not hold has probability zero under $\mathbb{P}$.

A particularly simple probability measure is the *delta* measure. For each $\omega \in \Omega$, we define a measure $\delta_\omega : \mathcal{H} \rightarrow [0, 1]$ such that $\delta_\omega(A) = \mathbf{1}_A(\omega)$, i.e. δ_ω outputs 1 if ω belongs to A and 0 otherwise. It can be trivially checked that δ_ω satisfies the two axioms of probability measures.

The following lemma collects some elementary properties of probability measures.

Lemma 2.3.1. *Let $(\Omega, \mathcal{H}, \mathbb{P})$ be a probability space.*

- (i) $\mathbb{P}(\emptyset) = 0$.
- (ii) If $A \subseteq B$, then $\mathbb{P}(A) \leq \mathbb{P}(B)$, and $\mathbb{P}(B \setminus A) = \mathbb{P}(B) - \mathbb{P}(A)$.
- (iii) If $A_1 \subseteq A_2 \subseteq \dots$, then $\mathbb{P}(\cup_{n \in \mathbb{N}} A_n) = \lim_{n \rightarrow \infty} \mathbb{P}(A_n)$.
- (iv) If $A_1 \supseteq A_2 \supseteq \dots$, then $\mathbb{P}(\cap_{n \in \mathbb{N}} A_n) = \lim_{n \rightarrow \infty} \mathbb{P}(A_n)$.
- (v) If $A_1, A_2, \dots \in \mathcal{H}$ (not necessarily disjoint), then $\mathbb{P}(\cup_{n \in \mathbb{N}} A_n) \leq \sum_{n \in \mathbb{N}} \mathbb{P}(A_n)$.

Proof. (i) For (i), apply σ -additivity to the disjoint sequence $\emptyset, \emptyset, \dots$. Since $\mathbb{P}(\emptyset) = \sum_{n \in \mathbb{N}} \mathbb{P}(\emptyset)$ and $\mathbb{P}(\emptyset) \in [0, 1]$, we must have $\mathbb{P}(\emptyset) = 0$.

(ii) Write $B = A \sqcup (B \setminus A)$. Then $\mathbb{P}(B) = \mathbb{P}(A) + \mathbb{P}(B \setminus A)$, which gives both claims.

(iii) Set $B_1 = A_1$ and $B_n = A_n \setminus A_{n-1}$ for $n \geq 2$. Then the B_n 's are pairwise disjoint and $\cup_{n \in \mathbb{N}} A_n = \sqcup_{n \in \mathbb{N}} B_n$. Hence

$$\mathbb{P}(\cup_{n \in \mathbb{N}} A_n) = \sum_{n \in \mathbb{N}} \mathbb{P}(B_n) = \lim_{m \rightarrow \infty} \sum_{n=1}^m \mathbb{P}(B_n) = \lim_{m \rightarrow \infty} \mathbb{P}(A_m).$$

- (iv) Apply (iii) to the increasing sequence $\Omega \setminus A_n$. Then $\Omega \setminus \bigcap_{n \in \mathbb{N}} A_n = \bigcup_{n \in \mathbb{N}} (\Omega \setminus A_n)$, so

$$1 - \mathbb{P}(\bigcap_{n \in \mathbb{N}} A_n) = \lim_{n \rightarrow \infty} (1 - \mathbb{P}(A_n)).$$

Rearranging gives the result.

- (v) Set $B_1 = A_1$, and for $n \geq 2$, $B_n = A_n \setminus \bigcup_{k=1}^{n-1} A_k$. Then these are disjoint with $\bigcup_{n \in \mathbb{N}} B_n = \bigcup_{n \in \mathbb{N}} A_n$ and $B_n \subseteq A_n$ for each $n \in \mathbb{N}$, so $\mathbb{P}(\bigcup_{n \in \mathbb{N}} A_n) = \mathbb{P}(\bigcup_{n \in \mathbb{N}} B_n) = \sum_{n \in \mathbb{N}} \mathbb{P}(B_n) \leq \sum_{n \in \mathbb{N}} \mathbb{P}(A_n)$, where we used σ -additivity and (ii). $\square$

For a probability measure $\mathbb{P}$ on a measurable space $(\Omega, \mathcal{H})$ and a sub- σ -algebra $\mathcal{F} \subseteq \mathcal{H}$, we denote the *restriction* of $\mathbb{P}$ on $\mathcal{F}$ by $\mathbb{P}|_{\mathcal{F}}$. If $(\Omega, \mathcal{H}) = \otimes_{t \in \mathbb{T}} (\Omega_t, \mathcal{E}_t)$ is a product space, then for any $\mathbb{S} \in \mathcal{P}(\mathbb{T})$, we abuse the notation slightly to write $\mathbb{P}|_{\mathbb{S}}$ for $\mathbb{P}|_{\mathcal{H}_{\mathbb{S}}}$. Sometimes, when it is obvious, we abuse the notation to simply write $\mathbb{P}$ without making the restriction explicit. Further, we have the *marginal* measure on $(\Omega_{\mathbb{S}}, \mathcal{E}_{\mathbb{S}})$ defined by $\mathbb{P} \circ \pi_{\mathbb{S}}^{-1}$.

In a probability space $(\Omega, \mathcal{H}, \mathbb{P})$, we say that a finite number of events $A_1, \dots, A_n \in \mathcal{H}$ are *independent* if, for every non-empty subset $I \subseteq \{1, \dots, n\}$,

$$\mathbb{P}(\bigcap_{i \in I} A_i) = \prod_{i \in I} \mathbb{P}(A_i).$$

We say that a finite number of sub- σ -algebras $\mathcal{H}_1, \dots, \mathcal{H}_n \subseteq \mathcal{H}$ are independent if all tuples of events $A_1 \in \mathcal{H}_1, \dots, A_n \in \mathcal{H}_n$ are independent. Finally, we say that an arbitrary set $\{\mathcal{H}_t\}_{t \in \mathbb{T}}$ of sub- σ -algebras of $\mathcal{H}$, where $\mathbb{T}$ can be uncountably infinite, are independent if every finite subset of them are independent.

If two probability measures agree on a π -system that generates a σ -algebra, then they agree on that σ -algebra.

Lemma 2.3.2. *Suppose that $\mathcal{C}$ is a π -system that generates $\mathcal{H}$, and $\mathbb{P}$ and $\mathbb{Q}$ are probability measures on $(\Omega, \mathcal{H})$ such that $\mathbb{P}(C) = \mathbb{Q}(C)$ for all $C \in \mathcal{C}$. Then $\mathbb{P}(A) = \mathbb{Q}(A)$ for all $A \in \mathcal{H}$.*

Proof. Define a collection $\mathcal{D} = \{A \in \mathcal{H} : \mathbb{P}(A) = \mathbb{Q}(A)\}$. Then $\mathcal{D}$ is a λ -system:

- $\mathbb{P}(\Omega) = \mathbb{Q}(\Omega)$, so $\Omega \in \mathcal{D}$;
- If $A, B \in \mathcal{D}$ with $A \supseteq B$, then by the additivity of probability measures,

$$\mathbb{P}(A \setminus B) = \mathbb{P}(A) - \mathbb{P}(B) = \mathbb{Q}(A) - \mathbb{Q}(B) = \mathbb{Q}(A \setminus B),$$

so $A \setminus B \in \mathcal{D}$.

- If $A_n \in \mathcal{D}$ for all $n \in \mathbb{N}$ with $A_1 \subseteq A_2 \subseteq \dots$, then Lemma 2.3.1(iii) tells us that

$$\mathbb{P}(\bigcup_{n \in \mathbb{N}} A_n) = \lim_{n \rightarrow \infty} \mathbb{P}(A_n) = \lim_{n \rightarrow \infty} \mathbb{Q}(A_n) = \mathbb{Q}(\bigcup_{n \in \mathbb{N}} A_n),$$

so $\bigcup_{n \in \mathbb{N}} A_n \in \mathcal{D}$.

Hence, by the monotone class theorem (Theorem 2.1.5), we have $\mathcal{D} = \mathcal{H}$, i.e. $\mathbb{P}(A) = \mathbb{Q}(A)$ for all $A \in \mathcal{H}$. $\square$

Let $(\Omega_1, \mathcal{H}_1)$ and $(\Omega_2, \mathcal{H}_2)$ be two measurable spaces. Then we say that a function $K : \Omega_1 \times \mathcal{H}_2 \rightarrow [0, 1]$ is a *transition probability kernel* from $(\Omega_1, \mathcal{H}_1)$ into $(\Omega_2, \mathcal{H}_2)$ if

- for each fixed $B \in \mathcal{H}_2$, the function $\omega_1 \mapsto K(\omega_1; B) : \Omega_1 \rightarrow [0, 1]$ is $\mathcal{H}_1$ -measurable; and
- for each fixed $\omega_1 \in \Omega_1$, the function $B \mapsto K(\omega_1; B) : \mathcal{H}_2 \rightarrow [0, 1]$ is a probability measure.

Transition probability kernels are the single most important object in the theory of causal spaces. We revisit them in Section 2.5. We present a result that greatly simplifies proofs that objects are transition probability kernels.

Lemma 2.3.3. *Let $(\Omega_1, \mathcal{H}_1)$ and $(\Omega_2, \mathcal{H}_2)$ be two measurable spaces, and $K : \Omega_1 \times \mathcal{H}_2 \rightarrow [0, 1]$ a function such that, for each $\omega \in \Omega_1$, the map $A \mapsto K(\omega; A) : \mathcal{H}_2 \rightarrow [0, 1]$ is a probability measure. Then the collection*

$$\mathcal{M} = \{A \in \mathcal{H}_2 : \omega \mapsto K(\omega; A) \text{ is } \mathcal{H}_1\text{-measurable}\} \subseteq \mathcal{H}_2$$

is a λ -system.

Hence, by applying the monotone class theorem (Theorem 2.1.5), if $\omega \mapsto K(\omega; A)$ is $\mathcal{H}_1$ -measurable for all A in a π -system that generates $\mathcal{H}_2$, then $\mathcal{M} = \mathcal{H}_2$, and K is a transition probability kernel.

Proof. Since $K(\omega; \Omega_2) = 1$ for all $\omega \in \Omega_1$, we have $\Omega_2 \in \mathcal{M}$. Now let $A, B \in \mathcal{M}$ with $A \subseteq B$. Then, for all $\omega \in \Omega_1$,

$$K(\omega; B \setminus A) = K(\omega; B) - K(\omega; A),$$

since $K(\omega; \cdot)$ is a probability measure. Since $\omega \mapsto K(\omega; A)$ and $\omega \mapsto K(\omega; B)$ are $\mathcal{H}_1$ -measurable, and subtractions of measurable functions are measurable by Lemma 2.2.2, it follows that $\omega \mapsto K(\omega; B \setminus A)$ is also $\mathcal{H}_1$ -measurable. Hence $B \setminus A \in \mathcal{M}$.

Finally, let $A_1 \subseteq A_2 \subseteq \dots$ with $A_n \in \mathcal{M}$, and set $A = \cup_{n \in \mathbb{N}} A_n$. Since $K(\omega; \cdot)$ is a probability measure for each ω , continuity from below (Lemma 2.3.1(iii)) gives

$$K(\omega; A) = \lim_{n \rightarrow \infty} K(\omega; A_n).$$

The right-hand side is the pointwise limit of $\mathcal{H}_1$ -measurable functions, and is therefore $\mathcal{H}_1$ -measurable by Theorem 2.2.6(i). Hence $A \in \mathcal{M}$. Therefore $\mathcal{M}$ is a λ -system. $\square$

Suppose that $(\Omega_1, \mathcal{H}_1)$ and $(\Omega_2, \mathcal{H}_2)$ are two measurable spaces, $\mathbb{P}$ a probability measure on $(\Omega_1, \mathcal{H}_1)$ and $f : \Omega_1 \rightarrow \Omega_2$ an $\mathcal{H}_1/\mathcal{H}_2$ -measurable function. Then we define the *image* of $\mathbb{P}$ under f as a map $\mathbb{P} \circ f^{-1} : \mathcal{H}_2 \rightarrow [0, 1]$ defined as

$$\mathbb{P} \circ f^{-1}(A) = \mathbb{P}(f^{-1}(A)).$$

It is easy to show that $\mathbb{P} \circ f^{-1}$ is a probability measure on $(\Omega_2, \mathcal{H}_2)$. Indeed, $f^{-1}(\Omega_2) = \Omega_1$ and $\mathbb{P}(\Omega_1) = 1$, so $\mathbb{P} \circ f^{-1}(\Omega_2) = 1$. Further, for a disjoint sequence $A_1, A_2, \dots \in \mathcal{H}_2$, the pre-images $f^{-1}(A_1), f^{-1}(A_2), \dots \in \mathcal{H}_1$ are also disjoint, and $f^{-1}(\sqcup_{n \in \mathbb{N}} A_n) = \sqcup_{n \in \mathbb{N}} f^{-1}(A_n)$. Hence by the σ -additivity of $\mathbb{P}$, we have

$$\mathbb{P} \circ f^{-1}(\sqcup_{n \in \mathbb{N}} A_n) = \mathbb{P}(\sqcup_{n \in \mathbb{N}} f^{-1}(A_n)) = \sum_{n \in \mathbb{N}} \mathbb{P}(f^{-1}(A_n)) = \sum_{n \in \mathbb{N}} \mathbb{P} \circ f^{-1}(A_n),$$

proving the σ -additivity of $\mathbb{P} \circ f^{-1}$.

2.4 Carathéodory extension theorem and product measures

Suppose that $\mathcal{S}$ is a semi-ring on Ω , which we recall from Section 2.1. The following is a classical extension result, simplified for the purposes of probability theory. The proof is still long, and should be skipped by most readers.

Theorem 2.4.1 (Carathéodory extension theorem). *Suppose that $\Omega \in \mathcal{S}$, and that a function $\mathbb{P}_0 : \mathcal{S} \rightarrow [0, 1]$ is a pre-measure, i.e. it satisfies*

- $\mathbb{P}_0(\emptyset) = 0$ and $\mathbb{P}_0(\Omega) = 1$;
- σ -additivity, in the sense that $\mathbb{P}_0(\sqcup_{n \in \mathbb{N}} A_n) = \sum_{n \in \mathbb{N}} \mathbb{P}_0(A_n)$ for all pairwise disjoint sequences $(A_n)_{n \in \mathbb{N}}$ in $\mathcal{S}$ such that $\sqcup_{n \in \mathbb{N}} A_n \in \mathcal{S}$ (note that a semi-ring is not necessarily closed under countable unions).

Then there exists a unique probability measure $\mathbb{P}$ on $\sigma(\mathcal{S})$, the σ -algebra generated by $\mathcal{S}$, such that $\mathbb{P}(A) = \mathbb{P}_0(A)$ for all $A \in \mathcal{S}$.

Proof. We proceed in several steps.

Step 1: Extend $\mathbb{P}_0$ from the semi-ring to the algebra of finite disjoint unions. Let $\mathcal{R}$ be the collection of all finite disjoint unions of elements of $\mathcal{S}$, that is,

$$\mathcal{R} = \{\sqcup_{i=1}^n S_i : n \in \mathbb{N}, S_i \in \mathcal{S}\}.$$

Since $\emptyset \in \mathcal{S}$, this collection contains $\emptyset$. We first show that $\mathcal{R}$ is closed under finite intersections and differences. Let $A = \sqcup_{i=1}^m S_i$ and $B = \sqcup_{j=1}^n T_j$, where $S_i, T_j \in \mathcal{S}$. Since $\mathcal{S}$ is closed under intersections, $A \cap B = \sqcup_{i=1}^m \sqcup_{j=1}^n (S_i \cap T_j)$ is again a finite disjoint union of elements of $\mathcal{S}$. Hence $A \cap B \in \mathcal{R}$.

Next, we show that $A \setminus B \in \mathcal{R}$. It is enough to show that $S \setminus B \in \mathcal{R}$ for each $S \in \mathcal{S}$, because then $A \setminus B = \sqcup_{i=1}^m (S_i \setminus B)$. Fix $S \in \mathcal{S}$. Since $B = \sqcup_{j=1}^n T_j$, we have $S \setminus B = (((S \setminus T_1) \setminus T_2) \cdots) \setminus T_n$. By the definition of a semi-ring, $S \setminus T_1$ can be written as a finite disjoint union of elements of $\mathcal{S}$. Applying the same argument successively to each of these pieces and to $T_2, \dots, T_n$, we obtain a finite disjoint decomposition of $S \setminus B$ into elements of $\mathcal{S}$. Hence $S \setminus B \in \mathcal{R}$, and therefore $A \setminus B \in \mathcal{R}$.

Since $\Omega \in \mathcal{S}$, we also have $\Omega \in \mathcal{R}$. Thus $\mathcal{R}$ is closed under complements relative to Ω . It follows that $\mathcal{R}$ is also closed under finite unions, because $A \cup B = A \sqcup (B \setminus A)$. Therefore $\mathcal{R}$ is an algebra of subsets of Ω .

Now define $\mathbb{P}_{\mathcal{R}} : \mathcal{R} \rightarrow [0, 1]$ as follows. If $A = \sqcup_{i=1}^m S_i$ with $S_i \in \mathcal{S}$, set

$$\mathbb{P}_{\mathcal{R}}(A) = \sum_{i=1}^m \mathbb{P}_0(S_i).$$

We need to show that this is well-defined. Suppose also that $A = \sqcup_{j=1}^n T_j$ with $T_j \in \mathcal{S}$. Then, for each i , we have $S_i = \sqcup_{j=1}^n (S_i \cap T_j)$, where each $S_i \cap T_j \in \mathcal{S}$. Since this is a finite disjoint union whose union is $S_i \in \mathcal{S}$, the finite version of the additivity assumption on $\mathbb{P}_0$ gives $\mathbb{P}_0(S_i) = \sum_{j=1}^n \mathbb{P}_0(S_i \cap T_j)$. Similarly, $\mathbb{P}_0(T_j) = \sum_{i=1}^m \mathbb{P}_0(S_i \cap T_j)$. Therefore

$$\sum_{i=1}^m \mathbb{P}_0(S_i) = \sum_{i=1}^m \sum_{j=1}^n \mathbb{P}_0(S_i \cap T_j) = \sum_{j=1}^n \sum_{i=1}^m \mathbb{P}_0(S_i \cap T_j) = \sum_{j=1}^n \mathbb{P}_0(T_j).$$

Thus $\mathbb{P}_{\mathcal{R}}$ is well-defined.

It is immediate from the definition that $\mathbb{P}_{\mathcal{R}}$ is finitely additive on $\mathcal{R}$. In particular, since $\Omega \in \mathcal{S}$, we have $\mathbb{P}_{\mathcal{R}}(\Omega) = \mathbb{P}_0(\Omega) = 1$, and by finite additivity, $\mathbb{P}_{\mathcal{R}}(A) \leq 1$ for every $A \in \mathcal{R}$. We now show that $\mathbb{P}_{\mathcal{R}}$ is σ -additive on $\mathcal{R}$. Let $A, A_1, A_2, \dots \in \mathcal{R}$ be such that the A_n are pairwise disjoint and $A = \sqcup_{n \in \mathbb{N}} A_n$. Choose finite disjoint decompositions $A = \sqcup_{i=1}^m S_i$ and $A_n = \sqcup_{k=1}^{m_n} S_{n,k}$, with all $S_i, S_{n,k} \in \mathcal{S}$. For each fixed i , we have $S_i = \sqcup_{n \in \mathbb{N}} \sqcup_{k=1}^{m_n} (S_i \cap S_{n,k})$, and all the sets $S_i \cap S_{n,k}$ belong to $\mathcal{S}$. Since their union is $S_i \in \mathcal{S}$, the σ -additivity assumption on $\mathbb{P}_0$ gives $\mathbb{P}_0(S_i) = \sum_{n \in \mathbb{N}} \sum_{k=1}^{m_n} \mathbb{P}_0(S_i \cap S_{n,k})$. Summing over $i = 1, \dots, m$, we obtain

$$\mathbb{P}_{\mathcal{R}}(A) = \sum_{i=1}^m \mathbb{P}_0(S_i) = \sum_{i=1}^m \sum_{n \in \mathbb{N}} \sum_{k=1}^{m_n} \mathbb{P}_0(S_i \cap S_{n,k}) = \sum_{n \in \mathbb{N}} \sum_{k=1}^{m_n} \sum_{i=1}^m \mathbb{P}_0(S_i \cap S_{n,k}).$$

For fixed n, k , since $S_{n,k} \subseteq A = \sqcup_{i=1}^m S_i$, we have $S_{n,k} = \sqcup_{i=1}^m (S_i \cap S_{n,k})$, and therefore $\mathbb{P}_0(S_{n,k}) = \sum_{i=1}^m \mathbb{P}_0(S_i \cap S_{n,k})$. Hence

$$\mathbb{P}_{\mathcal{R}}(A) = \sum_{n \in \mathbb{N}} \sum_{k=1}^{m_n} \mathbb{P}_0(S_{n,k}) = \sum_{n \in \mathbb{N}} \mathbb{P}_{\mathcal{R}}(A_n).$$

Thus $\mathbb{P}_{\mathcal{R}}$ is a probability pre-measure on the algebra $\mathcal{R}$.

Step 2: Define an outer measure. For every subset $A \subseteq \Omega$, define

$$\mathbb{P}^*(A) = \inf \left\{ \sum_{n \in \mathbb{N}} \mathbb{P}_{\mathcal{R}}(R_n) : A \subseteq \cup_{n \in \mathbb{N}} R_n, R_n \in \mathcal{R} \right\}.$$

This infimum is taken over a non-empty collection, because $A \subseteq \Omega$ and $\Omega \in \mathcal{R}$. We show that $\mathbb{P}^*$ is an outer measure, meaning that $\mathbb{P}^*(\emptyset) = 0$, that it is monotone, and that it is σ -subadditive.

First, since $\emptyset \in \mathcal{R}$ and $\mathbb{P}_{\mathcal{R}}(\emptyset) = 0$, we have $\mathbb{P}^*(\emptyset) = 0$. Next, if $A \subseteq B \subseteq \Omega$, then every countable cover of B by sets in $\mathcal{R}$ is also a countable cover of A . Hence $\mathbb{P}^*(A) \leq \mathbb{P}^*(B)$, and so $\mathbb{P}^*$ is monotone. Finally, let $A_1, A_2, \dots \subseteq \Omega$. We show that

$$\mathbb{P}^*\left(\bigcup_{n \in \mathbb{N}} A_n\right) \leq \sum_{n \in \mathbb{N}} \mathbb{P}^*(A_n).$$

Fix $\varepsilon > 0$. For each $n \in \mathbb{N}$, choose sets $R_{n,k} \in \mathcal{R}$, $k \in \mathbb{N}$, such that $A_n \subseteq \bigcup_{k \in \mathbb{N}} R_{n,k}$ and $\sum_{k \in \mathbb{N}} \mathbb{P}_{\mathcal{R}}(R_{n,k}) \leq \mathbb{P}^*(A_n) + \frac{\varepsilon}{2^n}$. Then the countable collection $(R_{n,k})_{n,k \in \mathbb{N}}$ covers $\bigcup_{n \in \mathbb{N}} A_n$. Therefore

$$\mathbb{P}^*\left(\bigcup_{n \in \mathbb{N}} A_n\right) \leq \sum_{n \in \mathbb{N}} \sum_{k \in \mathbb{N}} \mathbb{P}_{\mathcal{R}}(R_{n,k}) \leq \sum_{n \in \mathbb{N}} \mathbb{P}^*(A_n) + \varepsilon.$$

Since $\varepsilon > 0$ was arbitrary, σ -subadditivity follows.

Step 3: $\mathbb{P}^*$ agrees with $\mathbb{P}_{\mathcal{R}}$ on $\mathcal{R}$. Let $A \in \mathcal{R}$. Since A covers itself, $\mathbb{P}^*(A) \leq \mathbb{P}_{\mathcal{R}}(A)$. For the reverse inequality, let $R_1, R_2, \dots \in \mathcal{R}$ be any countable cover of A , so that $A \subseteq \bigcup_{n \in \mathbb{N}} R_n$. Define $B_1 = A \cap R_1$, and, for $n \geq 2$, $B_n = A \cap (R_n \setminus \bigcup_{k=1}^{n-1} R_k)$. Since $\mathcal{R}$ is closed under finite unions, intersections and differences, each $B_n \in \mathcal{R}$. The sets B_n are pairwise disjoint, and because the R_n cover A , $A = \bigsqcup_{n \in \mathbb{N}} B_n$. By σ -additivity of $\mathbb{P}_{\mathcal{R}}$ on $\mathcal{R}$, we have $\mathbb{P}_{\mathcal{R}}(A) = \sum_{n \in \mathbb{N}} \mathbb{P}_{\mathcal{R}}(B_n)$. Since $B_n \subseteq R_n$, finite additivity of $\mathbb{P}_{\mathcal{R}}$ gives $\mathbb{P}_{\mathcal{R}}(B_n) \leq \mathbb{P}_{\mathcal{R}}(R_n)$. Therefore $\mathbb{P}_{\mathcal{R}}(A) \leq \sum_{n \in \mathbb{N}} \mathbb{P}_{\mathcal{R}}(R_n)$. Taking the infimum over all countable covers of A by sets in $\mathcal{R}$, we obtain $\mathbb{P}_{\mathcal{R}}(A) \leq \mathbb{P}^*(A)$. Thus $\mathbb{P}^*(A) = \mathbb{P}_{\mathcal{R}}(A)$ for every $A \in \mathcal{R}$.

Step 4: Carathéodory measurable sets form a σ -algebra. We say that a set $A \subseteq \Omega$ is $\mathbb{P}^*$ -measurable if, for every $C \subseteq \Omega$,

$$\mathbb{P}^*(C) = \mathbb{P}^*(C \cap A) + \mathbb{P}^*(C \setminus A).$$

Let $\mathcal{M}$ be the collection of all $\mathbb{P}^*$ -measurable sets. We show that $\mathcal{M}$ is a σ -algebra. First, $\Omega \in \mathcal{M}$, since, for any $C \subseteq \Omega$,

$$\mathbb{P}^*(C) = \mathbb{P}^*(C \cap \Omega) + \mathbb{P}^*(C \setminus \Omega) = \mathbb{P}^*(C) + \mathbb{P}^*(\emptyset).$$

Also, if $A \in \mathcal{M}$, then $\Omega \setminus A \in \mathcal{M}$, because $C \cap (\Omega \setminus A) = C \setminus A$ and $C \setminus (\Omega \setminus A) = C \cap A$. Thus $\mathcal{M}$ is closed under complements.

Next, let $A, B \in \mathcal{M}$. We show that $A \cup B \in \mathcal{M}$. By σ -subadditivity, it is always true that, for any $C \subseteq \Omega$,

$$\mathbb{P}^*(C) \leq \mathbb{P}^*(C \cap (A \cup B)) + \mathbb{P}^*(C \setminus (A \cup B)).$$

Thus it suffices to prove the reverse inequality. Since $A \in \mathcal{M}$,

$$\mathbb{P}^*(C) = \mathbb{P}^*(C \cap A) + \mathbb{P}^*(C \setminus A).$$

Applying the measurability of B to both $C \cap A$ and $C \setminus A$, we get

$$\mathbb{P}^*(C) = \mathbb{P}^*(C \cap A \cap B) + \mathbb{P}^*((C \cap A) \setminus B) + \mathbb{P}^*((C \setminus A) \cap B) + \mathbb{P}^*((C \setminus A) \setminus B).$$

The first three sets on the right-hand side partition $C \cap (A \cup B)$, while the last set is exactly $C \setminus (A \cup B)$. By subadditivity,

$$\mathbb{P}^*(C \cap (A \cup B)) \leq \mathbb{P}^*(C \cap A \cap B) + \mathbb{P}^*((C \cap A) \setminus B) + \mathbb{P}^*((C \setminus A) \cap B).$$

Hence

$$\mathbb{P}^*(C) \geq \mathbb{P}^*(C \cap (A \cup B)) + \mathbb{P}^*(C \setminus (A \cup B)).$$

Therefore $A \cup B \in \mathcal{M}$. So $\mathcal{M}$ is closed under finite unions. Now let $A_1, A_2, \dots \in \mathcal{M}$ be pairwise disjoint, and set $A = \sqcup_{n \in \mathbb{N}} A_n$. For $m \in \mathbb{N}$, define $B_m = \sqcup_{n=1}^m A_n$. Since $\mathcal{M}$ is closed under finite unions, $B_m \in \mathcal{M}$. Repeatedly applying $\mathbb{P}^*$ -measurability gives, for every $C \subseteq \Omega$, $\mathbb{P}^*(C) = \sum_{n=1}^m \mathbb{P}^*(C \cap A_n) + \mathbb{P}^*(C \setminus B_m)$. Since $C \setminus B_m \supseteq C \setminus A$, monotonicity gives $\mathbb{P}^*(C \setminus B_m) \geq \mathbb{P}^*(C \setminus A)$. Therefore $\mathbb{P}^*(C) \geq \sum_{n=1}^m \mathbb{P}^*(C \cap A_n) + \mathbb{P}^*(C \setminus A)$. Letting $m \rightarrow \infty$, we obtain $\mathbb{P}^*(C) \geq \sum_{n \in \mathbb{N}} \mathbb{P}^*(C \cap A_n) + \mathbb{P}^*(C \setminus A)$. By σ -subadditivity, $\mathbb{P}^*(C \cap A) \leq \sum_{n \in \mathbb{N}} \mathbb{P}^*(C \cap A_n)$. Hence

$$\mathbb{P}^*(C) \geq \mathbb{P}^*(C \cap A) + \mathbb{P}^*(C \setminus A).$$

The reverse inequality follows from subadditivity, so $A \in \mathcal{M}$. Thus $\mathcal{M}$ is closed under countable disjoint unions. Finally, if $A_1, A_2, \dots \in \mathcal{M}$ are arbitrary, define the disjoint sets $B_1 = A_1$, and for $n \geq 2$, $B_n = A_n \setminus \cup_{k=1}^{n-1} A_k$. Since $\mathcal{M}$ is closed under complements and finite unions, each $B_n \in \mathcal{M}$. Also, $\cup_{n \in \mathbb{N}} A_n = \sqcup_{n \in \mathbb{N}} B_n$. Therefore $\mathcal{M}$ is closed under countable unions. Hence $\mathcal{M}$ is a σ -algebra.

Moreover, if $A_1, A_2, \dots \in \mathcal{M}$ are pairwise disjoint, then

$$\mathbb{P}^*(\sqcup_{n \in \mathbb{N}} A_n) = \sum_{n \in \mathbb{N}} \mathbb{P}^*(A_n).$$

Indeed, the inequality $\leq$ follows from σ -subadditivity. For the reverse inequality, finite additivity on $\mathcal{M}$ (which follows by letting $C = B_m$ above) gives $\mathbb{P}^*(\sqcup_{n \in \mathbb{N}} A_n) \geq \mathbb{P}^*(\sqcup_{n=1}^m A_n) = \sum_{n=1}^m \mathbb{P}^*(A_n)$ for every $m \in \mathbb{N}$. Letting $m \rightarrow \infty$ gives the claim.

Step 5: Every set in $\mathcal{R}$ is $\mathbb{P}^*$ -measurable. Let $A \in \mathcal{R}$. We show that $A \in \mathcal{M}$. Fix $C \subseteq \Omega$. By subadditivity, $\mathbb{P}^*(C) \leq \mathbb{P}^*(C \cap A) + \mathbb{P}^*(C \setminus A)$. Thus it remains to prove the reverse inequality.

Let $R_1, R_2, \dots \in \mathcal{R}$ be any countable cover of C . Then $(R_n \cap A)_{n \in \mathbb{N}}$ covers $C \cap A$, and $(R_n \setminus A)_{n \in \mathbb{N}}$ covers $C \setminus A$. Since $\mathcal{R}$ is closed under intersections and differences, all these sets belong to $\mathcal{R}$. Hence

$$\begin{aligned} \mathbb{P}^*(C \cap A) + \mathbb{P}^*(C \setminus A) &\leq \sum_{n \in \mathbb{N}} \mathbb{P}_{\mathcal{R}}(R_n \cap A) + \sum_{n \in \mathbb{N}} \mathbb{P}_{\mathcal{R}}(R_n \setminus A) \\ &= \sum_{n \in \mathbb{N}} [\mathbb{P}_{\mathcal{R}}(R_n \cap A) + \mathbb{P}_{\mathcal{R}}(R_n \setminus A)] \end{aligned}$$

$$= \sum_{n \in \mathbb{N}} \mathbb{P}_{\mathcal{R}}(R_n),$$

where the last equality follows from finite additivity of $\mathbb{P}_{\mathcal{R}}$, since $R_n = (R_n \cap A) \sqcup (R_n \setminus A)$. Taking the infimum over all countable covers of C by sets in $\mathcal{R}$, we get $\mathbb{P}^*(C \cap A) + \mathbb{P}^*(C \setminus A) \leq \mathbb{P}^*(C)$. Therefore $A \in \mathcal{M}$. Thus $\mathcal{R} \subseteq \mathcal{M}$.

Since $\mathcal{S} \subseteq \mathcal{R}$, we have $\sigma(\mathcal{S}) \subseteq \sigma(\mathcal{R})$. Conversely, every set in $\mathcal{R}$ is a finite union of elements of $\mathcal{S}$, so $\mathcal{R} \subseteq \sigma(\mathcal{S})$, and hence $\sigma(\mathcal{R}) = \sigma(\mathcal{S})$. Because $\mathcal{M}$ is a σ -algebra containing $\mathcal{R}$, it follows that $\sigma(\mathcal{S}) = \sigma(\mathcal{R}) \subseteq \mathcal{M}$.

Step 6: Define the extension. Define $\mathbb{P}$ on $\sigma(\mathcal{S})$ as the restriction of $\mathbb{P}^*$ to $\sigma(\mathcal{S})$. Since $\sigma(\mathcal{S}) \subseteq \mathcal{M}$, the σ -additivity shown in Step 4 implies that $\mathbb{P}$ is σ -additive on $\sigma(\mathcal{S})$. Also, $\mathbb{P}(\Omega) = \mathbb{P}^*(\Omega) = \mathbb{P}_{\mathcal{R}}(\Omega) = \mathbb{P}_0(\Omega) = 1$. Thus $\mathbb{P}$ is a probability measure on $\sigma(\mathcal{S})$. Finally, if $S \in \mathcal{S}$, then $S \in \mathcal{R}$, and by Step 3, $\mathbb{P}(S) = \mathbb{P}^*(S) = \mathbb{P}_{\mathcal{R}}(S) = \mathbb{P}_0(S)$. Thus $\mathbb{P}$ extends $\mathbb{P}_0$.

Step 7: Uniqueness. Suppose that $\mathbb{Q}$ is another probability measure on $\sigma(\mathcal{S})$ such that $\mathbb{Q}(S) = \mathbb{P}_0(S)$ for all $S \in \mathcal{S}$. We show that $\mathbb{Q} = \mathbb{P}$. Define

$$\mathcal{D} = \{A \in \sigma(\mathcal{S}) : \mathbb{P}(A) = \mathbb{Q}(A)\}.$$

We claim that $\mathcal{D}$ is a λ -system. First, since both $\mathbb{P}$ and $\mathbb{Q}$ are probability measures, we have $\mathbb{P}(\Omega) = 1 = \mathbb{Q}(\Omega)$, so $\Omega \in \mathcal{D}$. Next, if $A, B \in \mathcal{D}$ and $A \supseteq B$, then

$$\mathbb{P}(A \setminus B) = \mathbb{P}(A) - \mathbb{P}(B) = \mathbb{Q}(A) - \mathbb{Q}(B) = \mathbb{Q}(A \setminus B),$$

so $A \setminus B \in \mathcal{D}$. Finally, suppose $A_1 \subseteq A_2 \subseteq \dots$ and $A_n \in \mathcal{D}$ for all n . Let $A = \cup_{n \in \mathbb{N}} A_n$. Then $A = A_1 \sqcup \sqcup_{n=2}^{\infty} (A_n \setminus A_{n-1})$. Since $\mathcal{D}$ is closed under differences of nested sets, each $A_n \setminus A_{n-1} \in \mathcal{D}$. Therefore, by σ -additivity of $\mathbb{P}$ and $\mathbb{Q}$,

$$\mathbb{P}(A) = \mathbb{P}(A_1) + \sum_{n=2}^{\infty} \mathbb{P}(A_n \setminus A_{n-1}) = \mathbb{Q}(A_1) + \sum_{n=2}^{\infty} \mathbb{Q}(A_n \setminus A_{n-1}) = \mathbb{Q}(A).$$

Hence $A \in \mathcal{D}$. Thus $\mathcal{D}$ is a λ -system.

Now $\mathcal{S} \subseteq \mathcal{D}$, because both $\mathbb{P}$ and $\mathbb{Q}$ agree with $\mathbb{P}_0$ on $\mathcal{S}$. Also, $\mathcal{S}$ is a π -system, since a semi-ring is closed under intersections. By the monotone class theorem (Theorem 2.1.5), $\sigma(\mathcal{S}) \subseteq \mathcal{D}$. Therefore $\mathbb{P}(A) = \mathbb{Q}(A)$ for every $A \in \sigma(\mathcal{S})$. Hence the extension is unique. □

The prototypical application of the Carathéodory extension theorem is to define a *product* measure on a product measurable space. First, we introduce the expectation of simple functions (Eq. (2.3)). Expectation of general measurable functions will be the content of the next section, but the construction of the product probability measure requires the expectation of simple functions, so we take a brief detour on them here.

Let $f = \sum_{i=1}^n a_i \mathbf{1}_{A_i}$ be a simple function. Then we define the *expectation* of f as

$$\mathbb{E}[f] = \int \mathbb{P}(d\omega) f(\omega) = \sum_{i=1}^n a_i \mathbb{P}(A_i). \quad (2.5)$$

Note that the representation of a simple function is not unique, so we need to show that the expectation is nevertheless well-defined.

Lemma 2.4.2. *The expectation of a simple function f (Eq. (2.5)) is independent of the representation of f .*

Proof. Suppose that $f = \sum_{i=1}^n a_i \mathbf{1}_{A_i} = \sum_{j=1}^m b_j \mathbf{1}_{B_j}$, where the latter is in canonical form, with $(B_j)_{j=1}^m$ a disjoint partition of Ω and distinct b_j , which we know exists by Lemma 2.2.5. Recall the sets A'_ϵ from the proof of Lemma 2.2.5. Write $I \subseteq \{0, 1\}^n$ for the set of all $\epsilon \in \{0, 1\}^n$ such that $A'_\epsilon \neq \emptyset$. Because f is constant in each A'_ϵ with $\epsilon \in I$ and since b_j are distinct, each non-empty A'_ϵ belongs to exactly one B_j ; we write $j(\epsilon)$ for the j such that $A'_\epsilon \subseteq B_j$. Then for each $\epsilon \in I$, we have $\sum_{i:\epsilon_i=1} a_i = b_{j(\epsilon)}$, and further, each B_j is disjointly partitioned by $\{A'_\epsilon : \epsilon \in I, j(\epsilon) = j\}$, such that, by the additivity of $\mathbb{P}$, we have $\mathbb{P}(B_j) = \sum_{\epsilon \in I, j(\epsilon)=j} \mathbb{P}(A'_\epsilon)$. Then

$$\begin{aligned} \sum_{i=1}^n a_i \mathbb{P}(A_i) &= \sum_{i=1}^n a_i \sum_{\epsilon \in I: \epsilon_i=1} \mathbb{P}(A'_\epsilon) = \sum_{\epsilon \in I} \mathbb{P}(A'_\epsilon) \sum_{i:\epsilon_i=1} a_i \\ &= \sum_{\epsilon \in I} \mathbb{P}(A'_\epsilon) b_{j(\epsilon)} = \sum_{j=1}^m b_j \sum_{\epsilon \in I, j(\epsilon)=j} \mathbb{P}(A'_\epsilon) = \sum_{j=1}^m b_j \mathbb{P}(B_j). \end{aligned}$$

Hence, $\mathbb{E}[f]$ is independent of representation. $\square$

In particular, we have $\mathbb{E}[\mathbf{1}_A] = \mathbb{P}(A)$, i.e. the expectation of an indicator function $\mathbf{1}_A$ coincides with the probability of A .

We establish some elementary properties of expectations of simple functions in the following lemma.

Lemma 2.4.3. (i) *If f and g are simple functions and $a, b \in \mathbb{R}$ real constants, then $\mathbb{E}[af + bg] = a\mathbb{E}[f] + b\mathbb{E}[g]$.*

(ii) *If $f \leq g$ are simple functions, then $\mathbb{E}[f] \leq \mathbb{E}[g]$. In particular, letting $f = 0$, positive simple functions g satisfy $\mathbb{E}[g] \geq 0$.*

(iii) *If $(f_n)_{n \in \mathbb{N}}$ are simple functions with $f_1 \geq f_2 \geq \dots$, and $f_n(\omega) \rightarrow 0$ for every $\omega \in \Omega$, then $\mathbb{E}[f_n] \rightarrow 0$.*

Proof. Let $f = \sum_{i=1}^n a_i \mathbf{1}_{A_i}$ and $g = \sum_{j=1}^m b_j \mathbf{1}_{B_j}$ in their canonical forms (Lemma 2.2.5).

(i) We note first that $af + bg$ is indeed a simple function, hence $\mathbb{E}[af + bg]$ is well-defined by Eq. (2.5). Then by the finite additivity of $\mathbb{P}$,

$$\mathbb{E}[af + bg] = \sum_{i,j} (aa_i + bb_j) \mathbb{P}(A_i \cap B_j)$$

$$\begin{aligned}
 &= a \sum_{i=1}^n a_i \sum_{j=1}^m \mathbb{P}(A_i \cap B_j) + b \sum_{j=1}^m b_j \sum_{i=1}^n \mathbb{P}(A_i \cap B_j) \\
 &= a \sum_{i=1}^n a_i \mathbb{P}(A_i) + b \sum_{j=1}^m b_j \mathbb{P}(B_j) \\
 &= a\mathbb{E}[f] + b\mathbb{E}[g].
 \end{aligned}$$

- (ii) Now suppose $f \leq g$. Then whenever $A_i \cap B_j \neq \emptyset$, for any $\omega \in A_i \cap B_j$, we must have $a_i = f(\omega) \leq g(\omega) = b_j$, and so

$$\mathbb{E}[f] = \sum_{i,j} a_i \mathbb{P}(A_i \cap B_j) \leq \sum_{i,j} b_j \mathbb{P}(A_i \cap B_j) = \mathbb{E}[g].$$

- (iii) Since $(f_n)_{n \in \mathbb{N}}$ is decreasing with $f_n(\omega) \rightarrow 0$ for every $\omega \in \Omega$, we have $f_n \geq 0$ for every $n \in \mathbb{N}$; and f_1 , taking only finitely many values, satisfies $f_1 \leq c$ for some $c \in \mathbb{R}_+$. Fix $\varepsilon > 0$. The sets $E_n = \{\omega : f_n(\omega) > \varepsilon\}$ lie in $\mathcal{H}$ and decrease to $\emptyset$, so $\mathbb{P}(E_n) \rightarrow 0$ by Lemma 2.3.1(i) and (iv). Since $0 \leq f_n \leq f_1 \leq c$, we have $f_n \leq \varepsilon + c\mathbf{1}_{E_n}$ pointwise, whence $\mathbb{E}[f_n] \leq \varepsilon + c\mathbb{P}(E_n)$ by (i) and (ii). Taking $n \rightarrow \infty$, and then letting $\varepsilon \downarrow 0$, gives $\mathbb{E}[f_n] \rightarrow 0$.

□

Now we construct the product probability measure on product measurable spaces.

Proposition 2.4.4 (Product probability measure). *Take a product measurable space $(\Omega, \mathcal{H}) = \otimes_{t \in \mathbb{T}} (\Omega_t, \mathcal{E}_t)$ from Section 2.1, and a probability measure $\mathbb{P}_t$ on $(\Omega_t, \mathcal{E}_t)$ for each $t \in \mathbb{T}$. Then there exists a unique probability measure $\mathbb{P}$ on $(\Omega, \mathcal{H})$ such that, for every measurable cylinder $\times_{t \in \mathbb{T}} A_t$,*

$$\mathbb{P}(\times_{t \in \mathbb{T}} A_t) = \prod_{t \in \mathbb{T}} \mathbb{P}_t(A_t).$$

We denote this measure by $\otimes_{t \in \mathbb{T}} \mathbb{P}_t$, and call it the product probability measure.

Proof. Let $\mathcal{S}$ be the collection of measurable cylinders in $\Omega = \times_{t \in \mathbb{T}} \Omega_t$, which is a semi-ring by Lemma 2.1.6. We first define a set function $\mathbb{P}_0$ on $\mathcal{S}$. If $C = \times_{t \in \mathbb{T}} A_t$, set

$$\mathbb{P}_0(C) = \prod_{t \in \mathbb{T}} \mathbb{P}_t(A_t).$$

This product is well-defined because all but finitely many factors are equal to 1. Empty cylinders have value 0. We will show that this cylinder set function extends to a pre-measure on the algebra $\mathcal{R}$ of finite disjoint unions of cylinders, as constructed in Step 1 of the proof of Theorem 2.4.1. The desired probability measure will then follow from the Carathéodory extension theorem.

Step 1: Finite additivity on cylinders. We first prove the following elementary finite-dimensional fact. Suppose $\mathbb{S} \subseteq \mathbb{T}$ is finite, and suppose that a measurable rectangle $R = \times_{t \in \mathbb{S}} A_t$ is written as a finite disjoint union of measurable rectangles, i.e. $R = \sqcup_{i=1}^n R_i$ with $R_i = \times_{t \in \mathbb{S}} A_{i,t}$. Then

$$\prod_{t \in \mathbb{S}} \mathbb{P}_t(A_t) = \sum_{i=1}^n \prod_{t \in \mathbb{S}} \mathbb{P}_t(A_{i,t}).$$

Discarding any empty R_i changes neither side, since an empty rectangle has an empty side and so contributes a zero product; we may therefore assume that every R_i is non-empty, which forces $A_{i,t} \subseteq A_t$ for all i and all $t \in \mathbb{S}$.

We argue by induction on $|\mathbb{S}|$. If $|\mathbb{S}| = 1$, then $R = A_t = \sqcup_{i=1}^n A_{i,t} = \sqcup_{i=1}^n R_i$ for the unique $t \in \mathbb{S}$, and the identity is the finite additivity of $\mathbb{P}_t$. Now let $\mathbb{S} = \{t_1, \dots, t_m\}$ with $m \geq 2$, and assume the claim for index sets of size $m-1$. Fix $\omega \in \Omega_{t_1}$. If $\omega \in A_{t_1}$, then taking sections at ω exhibits $\times_{j=2}^m A_{t_j}$ as the disjoint union of the rectangles $\times_{j=2}^m A_{i,t_j}$ over those i with $\omega \in A_{i,t_1}$, so the induction hypothesis applies; and if $\omega \notin A_{t_1}$, then $\omega \notin A_{i,t_1}$ for every i . In either case,

$$\mathbf{1}_{A_{t_1}}(\omega) \prod_{j=2}^m \mathbb{P}_{t_j}(A_{t_j}) = \sum_{i=1}^n \mathbf{1}_{A_{i,t_1}}(\omega) \prod_{j=2}^m \mathbb{P}_{t_j}(A_{i,t_j}).$$

Both sides of this equation are simple functions on $(\Omega_{t_1}, \mathcal{E}_{t_1})$. Taking expectations with respect to $\mathbb{P}_{t_1}$ (Eq. (2.5)) and applying Lemma 2.4.3(i) repeatedly therefore gives the desired identity.

Finally, let a measurable cylinder $C = \times_{t \in \mathbb{T}} A_t$ be written as a finite disjoint union $C = \sqcup_{i=1}^n C_i$ of measurable cylinders $C_i = \times_{t \in \mathbb{T}} A_{i,t}$, and let $\mathbb{S} \subseteq \mathbb{T}$ be the finite set of coordinates at which C or some C_i differs from the full coordinate space. Since $\pi_{\mathbb{S}}$ is surjective and C and the C_i are the $\pi_{\mathbb{S}}$ -preimages of the corresponding rectangles in $\Omega_{\mathbb{S}}$, those rectangles form a disjoint decomposition of $\times_{t \in \mathbb{S}} A_t$. As all coordinates outside $\mathbb{S}$ contribute factors of 1, the identity above gives $\mathbb{P}_0(C) = \sum_{i=1}^n \mathbb{P}_0(C_i)$.

Step 2: Extend $\mathbb{P}_0$ to the cylinder algebra. Let $\mathcal{R}$ be the collection of finite disjoint unions of measurable cylinders:

$$\mathcal{R} = \left\{ \sqcup_{i=1}^n C_i : n \in \mathbb{N}, C_i \in \mathcal{S} \right\}.$$

Since $\mathcal{S}$ is a semi-ring, $\mathcal{R}$ is an algebra of subsets of Ω , by the reasoning given in Step 1 of the proof of Theorem 2.4.1. In particular, it contains Ω , and is closed under finite unions, finite intersections and complements.

Define $\mathbb{P}_{\mathcal{R}} : \mathcal{R} \rightarrow [0, 1]$ as in Step 1 of the proof of Theorem 2.4.1, i.e. $\mathbb{P}_{\mathcal{R}}(\sqcup_{i=1}^n C_i) = \sum_{i=1}^n \mathbb{P}_0(C_i)$, which we already showed to be well-defined. It is also finitely additive by construction, and $\mathbb{P}_{\mathcal{R}}(\emptyset) = 0$ and $\mathbb{P}_{\mathcal{R}}(\Omega) = 1$.

Step 3: σ -additivity with countable index set. Suppose in this step that $\mathbb{T}$ is countable, so that, relabelling, we may take $\mathbb{T} = \mathbb{N}$ or $\mathbb{T} = \{1, \dots, m\}$ for some $m \in \mathbb{N}$. In the latter case we adopt the convention that the empty product $\times_{k>n} \Omega_k$ for $n \geq m$ is a one-point space carrying the trivial probability measure, so that the argument below applies verbatim in both cases. We show that $\mathbb{P}_{\mathcal{R}}$ is σ -additive on $\mathcal{R}$.

It suffices to prove continuity from above, i.e. if $B_m \in \mathcal{R}$ with $B_1 \supseteq B_2 \supseteq \dots$ and $\cap_{m \in \mathbb{N}} B_m = \emptyset$, then $\mathbb{P}_{\mathcal{R}}(B_m) \rightarrow 0$. Indeed, if $A, A_1, A_2, \dots \in \mathcal{R}$, the A_n are pairwise disjoint, and $A = \sqcup_{n \in \mathbb{N}} A_n$, then $B_m = A \setminus \sqcup_{n=1}^m A_n$ lies in $\mathcal{R}$, since $\mathcal{R}$ is an algebra, and decreases to $\emptyset$, since every point of A lies in some A_n . As finite additivity gives $\mathbb{P}_{\mathcal{R}}(B_m) = \mathbb{P}_{\mathcal{R}}(A) - \sum_{n=1}^m \mathbb{P}_{\mathcal{R}}(A_n)$, continuity from above implies

$$\mathbb{P}_{\mathcal{R}}(A) = \sum_{n \in \mathbb{N}} \mathbb{P}_{\mathcal{R}}(A_n).$$

For $m \geq 0$, write $\Omega^{(m)} = \times_{k>m} \Omega_k$ for the tail product, let $\mathcal{R}^{(m)}$ be the algebra of finite disjoint unions of measurable cylinders in $\Omega^{(m)}$, and let $\mathbb{P}_{\mathcal{R}}^{(m)}$ be the corresponding measure of the cylinder, which is well-defined and finitely additive by Steps 1 and 2 applied to $\otimes_{k>m} (\Omega_k, \mathcal{E}_k)$; thus $\mathcal{R}^{(0)} = \mathcal{R}$ and $\mathbb{P}_{\mathcal{R}}^{(0)} = \mathbb{P}_{\mathcal{R}}$. For $R \in \mathcal{R}^{(m)}$ and $\eta \in \Omega_{m+1}$, write R_{η} for the section of R obtained by fixing the $(m+1)$ -st coordinate to be η ; since the section of a measurable cylinder is either a measurable cylinder or empty, and sectioning preserves disjoint unions, $R_{\eta} \in \mathcal{R}^{(m+1)}$. We write $R_{\omega_1, \dots, \omega_m} \in \mathcal{R}^{(m)}$ for the iterated section of $R \in \mathcal{R}$ fixing the first m coordinates, this being R itself when $m = 0$.

We shall use the following section identity: for $R \in \mathcal{R}^{(m)}$, the map $\eta \mapsto \mathbb{P}_{\mathcal{R}}^{(m+1)}(R_{\eta})$ is a simple function on $(\Omega_{m+1}, \mathcal{E}_{m+1}, \mathbb{P}_{m+1})$ with values in $[0, 1]$, and

$$\mathbb{P}_{\mathcal{R}}^{(m)}(R) = \int \mathbb{P}_{m+1}(d\eta) \mathbb{P}_{\mathcal{R}}^{(m+1)}(R_{\eta}).$$

If $R = \times_{k>m} A_k$, we have $\mathbb{P}_{\mathcal{R}}^{(m+1)}(R_{\eta}) = \mathbf{1}_{A_{m+1}}(\eta) \prod_{k>m+1} \mathbb{P}_k(A_k)$, whose expectation under $\mathbb{P}_{m+1}$ is $\prod_{k>m} \mathbb{P}_k(A_k) = \mathbb{P}_{\mathcal{R}}^{(m)}(R)$. For $R = \sqcup_{i=1}^p C_i$ we have $R_{\eta} = \sqcup_{i=1}^p (C_i)_{\eta}$, so the additivity of $\mathbb{P}_{\mathcal{R}}^{(m+1)}$ means $\mathbb{P}_{\mathcal{R}}^{(m+1)}(R_{\eta})$ is a simple function of η ; the identity then follows by applying Lemma 2.4.3(i) repeatedly and the single-cylinder case to each C_i .

Now suppose, for a contradiction, that $R_1 \supseteq R_2 \supseteq \dots$ lie in $\mathcal{R}$ with $\cap_{n \in \mathbb{N}} R_n = \emptyset$, but $a_0 := \lim_{n \rightarrow \infty} \mathbb{P}_{\mathcal{R}}(R_n) > 0$; the limit exists because $\mathbb{P}_{\mathcal{R}}$, being finitely additive, is monotone. We construct a point $\omega = (\omega_1, \omega_2, \dots) \in \Omega$ lying in every R_n , contradicting the assumption that the intersection is empty.

Suppose that $\omega_1, \dots, \omega_m$ have been chosen so that

$$a_m := \lim_{n \rightarrow \infty} \mathbb{P}_{\mathcal{R}}^{(m)}((R_n)_{\omega_1, \dots, \omega_m}) > 0,$$

the limit again existing by monotonicity; this holds for $m = 0$ by assumption. For $\eta \in \Omega_{m+1}$, set $f_n(\eta) = \mathbb{P}_{\mathcal{R}}^{(m+1)}((R_n)_{\omega_1, \dots, \omega_m, \eta})$. By the section identity applied to $(R_n)_{\omega_1, \dots, \omega_m} \in \mathcal{R}^{(m)}$, each f_n is a simple function with values in $[0, 1]$, and

$$\int \mathbb{P}_{m+1}(d\eta) f_n(\eta) = \mathbb{P}_{\mathcal{R}}^{(m)}((R_n)_{\omega_1, \dots, \omega_m}) \rightarrow a_m > 0, \quad \text{as } n \rightarrow \infty.$$

Moreover $f_1 \geq f_2 \geq \dots$, since $R_1 \supseteq R_2 \supseteq \dots$. Hence Lemma 2.4.3(iii) rules out $f_n(\eta) \rightarrow 0$ for every $\eta \in \Omega_{m+1}$, so there exists $\omega_{m+1} \in \Omega_{m+1}$ with

$$a_{m+1} := \lim_{n \rightarrow \infty} \mathbb{P}_{\mathcal{R}}^{(m+1)}((R_n)_{\omega_1, \dots, \omega_m, \omega_{m+1}}) > 0.$$

Continuing inductively, we obtain $\omega = (\omega_1, \omega_2, \dots) \in \Omega$.

We now show that $\omega \in R_k$ for every fixed $k \in \mathbb{N}$. Since $R_k \in \mathcal{R}$, membership of R_k is determined by the first m coordinates for some $m \in \mathbb{N}$. For every $n \geq k$ we have $(R_n)_{\omega_1, \dots, \omega_m} \subseteq (R_k)_{\omega_1, \dots, \omega_m}$, so monotonicity of $\mathbb{P}_{\mathcal{R}}^{(m)}$ and letting $n \rightarrow \infty$ give $\mathbb{P}_{\mathcal{R}}^{(m)}((R_k)_{\omega_1, \dots, \omega_m}) \geq a_m > 0$. But this section is either $\Omega^{(m)}$ or $\emptyset$, so it must be $\Omega^{(m)}$, and therefore $\omega \in R_k$.

Since k was arbitrary, $\omega \in \bigcap_{k \in \mathbb{N}} R_k$, a contradiction. Thus $\mathbb{P}_{\mathcal{R}}(R_n) \rightarrow 0$, and $\mathbb{P}_{\mathcal{R}}$ is σ -additive on $\mathcal{R}$ whenever $\mathbb{T}$ is countable.

Step 4: Reduction from arbitrary products to countable products. Now let $\mathbb{T}$ be arbitrary. Suppose $A, A_1, A_2, \dots \in \mathcal{R}$, with $A = \sqcup_{n \in \mathbb{N}} A_n$. Each of these sets is a finite disjoint union of measurable cylinders. Let $\mathbb{S} \subseteq \mathbb{T}$ be the union of all coordinates appearing in fixed finite cylinder decompositions of $A, A_1, A_2, \dots$. Then $\mathbb{S}$ is countable.

All the sets $A, A_1, A_2, \dots$ depend only on the coordinates in $\mathbb{S}$. Hence there exist corresponding sets $\tilde{A}, \tilde{A}_1, \tilde{A}_2, \dots \in \mathcal{R}_{\mathbb{S}}$ in the countable subproduct $\Omega_{\mathbb{S}} = \times_{s \in \mathbb{S}} \Omega_s$ such that $A = \pi_{\mathbb{S}}^{-1}(\tilde{A})$ and $A_n = \pi_{\mathbb{S}}^{-1}(\tilde{A}_n)$, where $\mathcal{R}_{\mathbb{S}}$ denotes the algebra of finite disjoint unions of measurable cylinders in the subproduct $\Omega_{\mathbb{S}}$. Since the original union is disjoint and equal to A , we also have $\tilde{A} = \sqcup_{n \in \mathbb{N}} \tilde{A}_n$. By the countable case, $\mathbb{P}_{\mathcal{R}, \mathbb{S}}(\tilde{A}) = \sum_{n \in \mathbb{N}} \mathbb{P}_{\mathcal{R}, \mathbb{S}}(\tilde{A}_n)$, where $\mathbb{P}_{\mathcal{R}, \mathbb{S}}$ is the set function on $\mathcal{R}_{\mathbb{S}}$ analogous to $\mathbb{P}_{\mathcal{R}}$ built from $(\mathbb{P}_t)_{t \in \mathbb{S}}$. The values agree with the original cylinder-algebra content, because all coordinates outside $\mathbb{S}$ contribute factors of 1. Therefore

$$\mathbb{P}_{\mathcal{R}}(A) = \sum_{n \in \mathbb{N}} \mathbb{P}_{\mathcal{R}}(A_n).$$

Thus $\mathbb{P}_{\mathcal{R}}$ is σ -additive on $\mathcal{R}$ for an arbitrary index set $\mathbb{T}$. Hence $\mathbb{P}_{\mathcal{R}}$ is a pre-measure on the algebra $\mathcal{R}$, which in particular is a semi-ring.

Step 5: Carathéodory. By the Carathéodory extension theorem (Theorem 2.4.1), there exists a unique probability measure $\mathbb{P}$ on $\sigma(\mathcal{R}) = \mathcal{H}$ such that $\mathbb{P}(A) = \mathbb{P}_{\mathcal{R}}(A)$ for all $A \in \mathcal{R}$. In particular, for every measurable cylinder $C = \times_{t \in \mathbb{T}} A_t$,

$$\mathbb{P}(C) = \mathbb{P}_0(C) = \prod_{t \in \mathbb{T}} \mathbb{P}_t(A_t).$$

If $\mathbb{Q}$ is another probability measure with the same values on measurable cylinders, then by finite additivity it has the same values as $\mathbb{P}_{\mathcal{R}}$ on every finite disjoint union of cylinders. Hence $\mathbb{Q}$ also agrees with $\mathbb{P}_{\mathcal{R}}$ on $\mathcal{R}$, and the uniqueness part of Theorem 2.4.1 gives $\mathbb{Q} = \mathbb{P}$ on $\mathcal{H}$.

□

If, instead, we have a probability measure $\mathbb{P}_t$ on $(\Omega, \mathcal{H}_t)$ for each $t \in \mathbb{T}$, then we have the image measures $\mathbb{P}_t \circ \pi_t^{-1}$ on $(\Omega_t, \mathcal{E}_t)$. Then following Proposition 2.4.4, we have a product probability measure $\otimes_{t \in \mathbb{T}} \mathbb{P}_t \circ \pi_t^{-1}$ on $(\Omega, \mathcal{H})$. By an abuse of notation, we suppress the canonical projections, and simply write $\otimes_{t \in \mathbb{T}} \mathbb{P}_t$.

2.5 Integration and expectation

Let us take a probability space $(\Omega, \mathcal{H}, \mathbb{P})$. In this section, we will define the expectation of $\overline{\mathbb{R}}_+$ -valued random variables, as well as $\overline{\mathbb{R}}$ -valued random variables satisfying an *integrability* condition.

Recall that, for simple functions f , the expectation $\mathbb{E}[f]$ was already defined in Section 2.4, in Eq. (2.5). Suppose now that $f : \Omega \rightarrow \overline{\mathbb{R}}_+$ is an arbitrary positive $\mathcal{H}$ -measurable function. Then by Theorem 2.2.6, we can write f as a pointwise limit of an increasing sequence of positive simple functions as $f(\omega) = \lim_{n \rightarrow \infty} d_n \circ f(\omega)$ for each $\omega \in \Omega$, where each d_n is recalled from Eq. (2.4). Then we define the *expectation* of f as

$$\mathbb{E}[f] = \int \mathbb{P}(d\omega) f(\omega) = \lim_{n \rightarrow \infty} \mathbb{E}[d_n \circ f]. \quad (2.6)$$

Note that, thanks to the monotonicity of $n \mapsto d_n$ and Lemma 2.4.3(ii), this limit exists, but can be infinite. Note also that, for any positive simple function $g = \sum_{i=1}^k a_i \mathbf{1}_{A_i}$, the values $d_m(a_i)$ converge to a_i as $m \rightarrow \infty$ for each i (since eventually $a_i < m$, and then $d_m(a_i) = \lfloor 2^m a_i \rfloor / 2^m$), so $\mathbb{E}[d_m \circ g] = \sum_{i=1}^k d_m(a_i) \mathbb{P}(A_i) \rightarrow \mathbb{E}[g]$; this is consistent with the simple-function definition of $\mathbb{E}[g]$ (Eq. (2.5)) and with the definition for positive random variables (Eq. (2.6)) obtained by applying Eq. (2.4) to g itself.

Each $d_n \circ f$ is a positive simple function, so by applying Lemma 2.4.3, we have the following positivity property of measurable functions $f : \Omega \rightarrow \overline{\mathbb{R}}_+$:

$$\mathbb{E}[f] = \lim_{n \rightarrow \infty} \mathbb{E}[d_n \circ f] \geq 0. \quad (2.7)$$

Further, the monotonicity property also follows easily: if $f, g : \Omega \rightarrow \overline{\mathbb{R}}_+$ are measurable functions with $f \leq g$, we have

$$\mathbb{E}[f] \leq \mathbb{E}[g]. \quad (2.8)$$

Indeed, the monotonicity of d_n gives $d_n \circ f \leq d_n \circ g$ for each $n \in \mathbb{N}$, and Lemma 2.4.3(ii) yields $\mathbb{E}[d_n \circ f] \leq \mathbb{E}[d_n \circ g]$. Taking $n \rightarrow \infty$, the definition of expectation for positive random variables gives $\mathbb{E}[f] = \lim_{n \rightarrow \infty} \mathbb{E}[d_n \circ f] \leq \lim_{n \rightarrow \infty} \mathbb{E}[d_n \circ g] = \mathbb{E}[g]$.

Finally, let $f : \Omega \rightarrow \overline{\mathbb{R}}$ be an arbitrary random variable. The *expectation* of f is defined as

$$\mathbb{E}[f] = \begin{cases} \mathbb{E}[f_+] - \mathbb{E}[f_-] & \text{if at least one of } \mathbb{E}[f_+] \text{ and } \mathbb{E}[f_-] \text{ is finite} \\ \text{undefined} & \text{otherwise.} \end{cases} \quad (2.9)$$

We say that a random variable f is *integrable* if $\mathbb{E}[f]$ is defined and is finite.

The following theorem is a key result that allows us to switch the order of integration and limit.

Theorem 2.5.1 (Monotone convergence theorem). *Let $(f_n)_{n \in \mathbb{N}}$ be an increasing sequence of random variables $f_n : \Omega \rightarrow \overline{\mathbb{R}}_+$. Then*

$$\mathbb{E} \left[\lim_{n \rightarrow \infty} f_n \right] = \lim_{n \rightarrow \infty} \mathbb{E}[f_n].$$

Proof. Note first that $\lim_{n \rightarrow \infty} f_n$ is well-defined, since (f_n) is increasing. Write $f = \lim_{n \rightarrow \infty} f_n$, so that $f : \Omega \rightarrow \overline{\mathbb{R}}_+$ is measurable by Theorem 2.2.6(i).

For one direction, since $f_n \leq f_{n+1} \leq f$ for each $n \in \mathbb{N}$, we have $d_m \circ f_n \leq d_m \circ f_{n+1} \leq d_m \circ f$ for each $m \in \mathbb{N}$ by the monotonicity of d_m . Applying Lemma 2.4.3(ii) gives $\mathbb{E}[d_m \circ f_n] \leq \mathbb{E}[d_m \circ f_{n+1}] \leq \mathbb{E}[d_m \circ f]$, and taking $m \rightarrow \infty$, the definition of expectation for positive random variables (Eq. (2.6)) yields $\mathbb{E}[f_n] \leq \mathbb{E}[f_{n+1}] \leq \mathbb{E}[f]$. Hence, $(\mathbb{E}[f_n])_{n \in \mathbb{N}}$ is an increasing sequence bounded above by $\mathbb{E}[f]$, so $\lim_{n \rightarrow \infty} \mathbb{E}[f_n]$ is well-defined and

$$\lim_{n \rightarrow \infty} \mathbb{E}[f_n] \leq \mathbb{E}[f].$$

For the reverse direction, we first claim that, for any positive simple function g with $g \leq f$,

$$\mathbb{E}[g] \leq \lim_{n \rightarrow \infty} \mathbb{E}[f_n]. \quad (*)$$

Write $g = \sum_{i=1}^k a_i \mathbf{1}_{A_i}$ in canonical form, with $a_i \in \mathbb{R}_+$. Fix $\alpha \in (0, 1)$, and define

$$E_n = \{\omega \in \Omega : f_n(\omega) \geq \alpha g(\omega)\} \in \mathcal{H}.$$

Since (f_n) is increasing, (E_n) is increasing too. Moreover, $\cup_{n \in \mathbb{N}} E_n = \Omega$: for any $\omega \in \Omega$, if $g(\omega) = 0$ then $\omega \in E_n$ for all n since $f_n(\omega) \geq 0 = \alpha g(\omega)$, whereas if $g(\omega) > 0$ then $\alpha g(\omega) < g(\omega) \leq f(\omega) = \lim_{n \rightarrow \infty} f_n(\omega)$, so $f_n(\omega) \geq \alpha g(\omega)$ eventually. Therefore $A_i \cap E_n \uparrow A_i$ as $n \rightarrow \infty$ for each i , and continuity of probability from below (Lemma 2.3.1(iii)) gives $\mathbb{P}(A_i \cap E_n) \rightarrow \mathbb{P}(A_i)$.

Now, the function $\alpha g \mathbf{1}_{E_n} = \sum_{i=1}^k \alpha a_i \mathbf{1}_{A_i \cap E_n}$ is a positive simple function that satisfies $\alpha g \mathbf{1}_{E_n} \leq f_n$ pointwise on Ω . By monotonicity of d_m , this gives $d_m \circ (\alpha g \mathbf{1}_{E_n}) \leq d_m \circ f_n$, and so by Lemma 2.4.3(ii), $\mathbb{E}[d_m \circ (\alpha g \mathbf{1}_{E_n})] \leq \mathbb{E}[d_m \circ f_n]$. Taking $m \rightarrow \infty$, the left side converges to $\mathbb{E}[\alpha g \mathbf{1}_{E_n}] = \sum_{i=1}^k \alpha a_i \mathbb{P}(A_i \cap E_n)$ since $\alpha g \mathbf{1}_{E_n}$ is simple, and the right side tends to $\mathbb{E}[f_n]$ by definition, so we obtain $\sum_{i=1}^k \alpha a_i \mathbb{P}(A_i \cap E_n) \leq \mathbb{E}[f_n]$. Taking $n \rightarrow \infty$, and using $\mathbb{P}(A_i \cap E_n) \rightarrow \mathbb{P}(A_i)$, we have $\alpha \mathbb{E}[g] = \alpha \sum_{i=1}^k a_i \mathbb{P}(A_i) \leq \lim_{n \rightarrow \infty} \mathbb{E}[f_n]$. Since this holds for all $\alpha \in (0, 1)$, letting $\alpha \uparrow 1$ yields $(*)$.

To conclude, let $g_m = d_m \circ f$ be the canonical approximating sequence given by Eq. (2.4), which is an increasing sequence of positive simple functions with $g_m \leq f$ and $\lim_{m \rightarrow \infty} g_m = f$ pointwise. Applying (*) to each g_m ,

$$\mathbb{E}[f] = \lim_{m \rightarrow \infty} \mathbb{E}[g_m] \leq \lim_{n \rightarrow \infty} \mathbb{E}[f_n],$$

where the first equality is the definition of $\mathbb{E}[f]$ (Eq. (2.6)). Combining with the previous direction gives $\mathbb{E}[f] = \lim_{n \rightarrow \infty} \mathbb{E}[f_n]$, as required. $\square$

The monotone convergence theorem allows us to obtain the linearity of expectations for general random variables, extending Lemma 2.4.3(i).

Lemma 2.5.2. (i) If $f, g : \Omega \rightarrow \overline{\mathbb{R}}_+$ are random variables and $a, b \geq 0$ are real numbers, then $\mathbb{E}[af + bg] = a\mathbb{E}[f] + b\mathbb{E}[g]$.

(ii) If $f, g : \Omega \rightarrow \overline{\mathbb{R}}$ are integrable random variables and $a, b \in \mathbb{R}$ are real numbers, then $\mathbb{E}[af + bg] = a\mathbb{E}[f] + b\mathbb{E}[g]$, provided that $af + bg$ is well-defined.

Proof. (i) The sequence $(ad_m \circ f + bd_m \circ g)_{m \in \mathbb{N}}$ is an increasing sequence of positive simple functions converging pointwise to $af + bg$ (Theorem 2.2.6). Applying the monotone convergence theorem (Theorem 2.5.1) to this sequence,

$$\mathbb{E}[af + bg] = \mathbb{E} \left[\lim_{m \rightarrow \infty} (ad_m \circ f + bd_m \circ g) \right] = \lim_{m \rightarrow \infty} \mathbb{E}[ad_m \circ f + bd_m \circ g].$$

By Lemma 2.4.3(i), $\mathbb{E}[ad_m \circ f + bd_m \circ g] = a\mathbb{E}[d_m \circ f] + b\mathbb{E}[d_m \circ g]$ for each m , and since $(a\mathbb{E}[d_m \circ f])_{m \in \mathbb{N}}$ and $(b\mathbb{E}[d_m \circ g])_{m \in \mathbb{N}}$ are increasing sequences in $[0, \infty]$, each having a well-defined limit (possibly equal to ∞), we may split the limit:

$$\mathbb{E}[af + bg] = \lim_{m \rightarrow \infty} a\mathbb{E}[d_m \circ f] + \lim_{m \rightarrow \infty} b\mathbb{E}[d_m \circ g] = a\mathbb{E}[f] + b\mathbb{E}[g].$$

(ii) Decompose $f = f_+ - f_-$, $g = g_+ - g_-$ and $f + g = (f + g)_+ - (f + g)_-$, where each of $f_{\pm}, g_{\pm}, (f + g)_{\pm}$ is a positive random variable. Rearranging the identity $f_+ - f_- + g_+ - g_- = (f + g)_+ - (f + g)_-$ so that both sides involve only positive random variables yields

$$(f + g)_+ + f_- + g_- = (f + g)_- + f_+ + g_+. \quad (*)$$

Each of these six positive random variables has finite expectation: indeed, we have $\mathbb{E}[f_{\pm}], \mathbb{E}[g_{\pm}] < \infty$ by the integrability of f, g , and the bound $(f + g)_{\pm} \leq |f| + |g| = f_+ + f_- + g_+ + g_-$ combined with (i) and Eq. (2.8) gives

$$\mathbb{E}[(f + g)_{\pm}] \leq \mathbb{E}[f_+] + \mathbb{E}[f_-] + \mathbb{E}[g_+] + \mathbb{E}[g_-] < \infty.$$

Applying (i) to both sides of (*) therefore yields

$$\mathbb{E}[(f + g)_+] + \mathbb{E}[f_-] + \mathbb{E}[g_-] = \mathbb{E}[(f + g)_-] + \mathbb{E}[f_+] + \mathbb{E}[g_+],$$

and rearranging—which is now legitimate, as all terms are finite—gives

$$\mathbb{E}[(f + g)_+] - \mathbb{E}[(f + g)_-] = (\mathbb{E}[f_+] - \mathbb{E}[f_-]) + (\mathbb{E}[g_+] - \mathbb{E}[g_-]),$$

i.e. $\mathbb{E}[f + g] = \mathbb{E}[f] + \mathbb{E}[g]$, as required.

Next we record homogeneity. For $a \geq 0$ we have $(af)_\pm = af_\pm$, whence $\mathbb{E}[af] = a\mathbb{E}[f]$; and $(-f)_\pm = f_\mp$ gives $\mathbb{E}[-f] = -\mathbb{E}[f]$. Combining the two cases yields $\mathbb{E}[af] = a\mathbb{E}[f]$ for every $a \in \mathbb{R}$. $\square$

We will also often need the following result.

Theorem 2.5.3. *Suppose that $f, g : \Omega \rightarrow \overline{\mathbb{R}}$ are random variables that are either both integrable or both take values in $\mathbb{R}_+$.*

- (i) *If $f(\omega) = g(\omega)$ for $\mathbb{P}$ -almost all $\omega \in \Omega$, then $\mathbb{E}[f] = \mathbb{E}[g]$.*
- (ii) *A kind of converse is that, if $\mathbb{E}[\mathbf{1}_A f] = \mathbb{E}[\mathbf{1}_A g]$ for all $A \in \mathcal{H}$, then $f(\omega) = g(\omega)$ for $\mathbb{P}$ -almost all $\omega \in \Omega$.*
- (iii) *The inequality version is that, if $\mathbb{E}[\mathbf{1}_A f] \leq \mathbb{E}[\mathbf{1}_A g]$ for all $A \in \mathcal{H}$, then $f(\omega) \leq g(\omega)$ for $\mathbb{P}$ -almost all $\omega \in \Omega$.*

Proof. (i) Suppose first that $f, g : \Omega \rightarrow \overline{\mathbb{R}}_+$. Let $A = \{f \neq g\} \in \mathcal{H}$, so that $\mathbb{P}(A) = 0$ by hypothesis. For each $m \in \mathbb{N}$, $d_m \circ f$ and $d_m \circ g$ are positive simple functions, and $d_m \circ f = d_m \circ g$ on $\Omega \setminus A$. Decomposing $d_m \circ f = (d_m \circ f)\mathbf{1}_{\Omega \setminus A} + (d_m \circ f)\mathbf{1}_A$ and applying Lemma 2.4.3(i),

$$\mathbb{E}[d_m \circ f] = \mathbb{E}[(d_m \circ f)\mathbf{1}_{\Omega \setminus A}] + \mathbb{E}[(d_m \circ f)\mathbf{1}_A].$$

Writing $d_m \circ f = \sum_k c_k \mathbf{1}_{A_k}$ with a disjoint partition $(A_k)_k$, $A_k \in \mathcal{H}$ of Ω , we have $(d_m \circ f)\mathbf{1}_A = \sum_k c_k \mathbf{1}_{A_k \cap A}$, and $\mathbb{P}(A_k \cap A) \leq \mathbb{P}(A) = 0$ for each k , so the second expectation vanishes. Hence $\mathbb{E}[d_m \circ f] = \mathbb{E}[(d_m \circ f)\mathbf{1}_{\Omega \setminus A}]$, and by the same argument $\mathbb{E}[d_m \circ g] = \mathbb{E}[(d_m \circ g)\mathbf{1}_{\Omega \setminus A}]$. Since $(d_m \circ f)\mathbf{1}_{\Omega \setminus A} = (d_m \circ g)\mathbf{1}_{\Omega \setminus A}$ pointwise on Ω , we conclude $\mathbb{E}[d_m \circ f] = \mathbb{E}[d_m \circ g]$, and taking $m \rightarrow \infty$ gives $\mathbb{E}[f] = \mathbb{E}[g]$.

For integrable $f, g : \Omega \rightarrow \overline{\mathbb{R}}$ with $f = g$ almost surely, we have $f_\pm = g_\pm$ almost surely as well. Applying the $\mathbb{R}_+$ case to $f_\pm$ and $g_\pm$ gives $\mathbb{E}[f_\pm] = \mathbb{E}[g_\pm]$, and so $\mathbb{E}[f] = \mathbb{E}[f_+] - \mathbb{E}[f_-] = \mathbb{E}[g_+] - \mathbb{E}[g_-] = \mathbb{E}[g]$.

- (ii) We show that $\mathbb{P}(f > g) = 0$; by symmetry, $\mathbb{P}(g > f) = 0$ as well, and so $f = g$ almost surely. For each $n \in \mathbb{N}$, let $A_n = \{f \geq g + \frac{1}{n}\} \cap \{g < \infty\} \in \mathcal{H}$, so that $\{f > g\} = \cup_{n \in \mathbb{N}} A_n$.

Suppose first that f, g are integrable. Then for each n , $\mathbf{1}_{A_n} f$ and $\mathbf{1}_{A_n} g$ are integrable, since $(\mathbf{1}_{A_n} f)_\pm = \mathbf{1}_{A_n} f_\pm \leq f_\pm$ gives $\mathbb{E}[(\mathbf{1}_{A_n} f)_\pm] \leq \mathbb{E}[f_\pm] < \infty$ by Eq. (2.8), and similarly for g . Applying Lemma 2.5.2 together with the hypothesis,

$$\mathbb{E}[\mathbf{1}_{A_n}(f - g)] = \mathbb{E}[\mathbf{1}_{A_n} f] - \mathbb{E}[\mathbf{1}_{A_n} g] = 0.$$

On A_n , $f - g \geq \frac{1}{n}$, so $\mathbf{1}_{A_n}(f - g)$ is a positive random variable with $\mathbf{1}_{A_n}(f - g) \geq \frac{1}{n}\mathbf{1}_{A_n}$ pointwise on Ω . By monotonicity (Eq. (2.8)),

$$0 = \mathbb{E}[\mathbf{1}_{A_n}(f - g)] \geq \frac{1}{n}\mathbb{E}[\mathbf{1}_{A_n}] = \frac{1}{n}\mathbb{P}(A_n),$$

hence $\mathbb{P}(A_n) = 0$. Therefore, by countable subadditivity of $\mathbb{P}$ (Lemma 2.3.1(v)), we have $\mathbb{P}(f > g) = \mathbb{P}(\cup_{n \in \mathbb{N}} A_n) \leq \sum_{n \in \mathbb{N}} \mathbb{P}(A_n) = 0$.

Suppose instead that $f, g : \Omega \rightarrow \overline{\mathbb{R}}_+$. For each $n, k \in \mathbb{N}$, let $A_{n,k} = A_n \cap \{g \leq k\} \in \mathcal{H}$, so that $A_n = \cup_{k \in \mathbb{N}} A_{n,k}$ by construction. Then $\mathbf{1}_{A_{n,k}}g \leq k$ is integrable, and the hypothesis gives $\mathbb{E}[\mathbf{1}_{A_{n,k}}f] = \mathbb{E}[\mathbf{1}_{A_{n,k}}g] \leq k < \infty$, so $\mathbf{1}_{A_{n,k}}f$ is integrable too. The same argument as above gives

$$0 = \mathbb{E}[\mathbf{1}_{A_{n,k}}f] - \mathbb{E}[\mathbf{1}_{A_{n,k}}g] = \mathbb{E}[\mathbf{1}_{A_{n,k}}(f - g)] \geq \frac{1}{n}\mathbb{P}(A_{n,k}),$$

so $\mathbb{P}(A_{n,k}) = 0$ for all $k \in \mathbb{N}$, and hence $\mathbb{P}(A_n) = \mathbb{P}(\cup_{k \in \mathbb{N}} A_{n,k}) \leq \sum_{k \in \mathbb{N}} \mathbb{P}(A_{n,k}) = 0$, by countable subadditivity (Lemma 2.3.1(v)). Therefore $\mathbb{P}(f > g) = \mathbb{P}(\cup_{n \in \mathbb{N}} A_n) \leq \sum_{n \in \mathbb{N}} \mathbb{P}(A_n) = 0$ by Lemma 2.3.1(v) again, as required.

- (iii) For $n, k \in \mathbb{N}$, the set $A_{n,k} = \{f \geq g + \frac{1}{n}\} \cap \{|g| \leq k\}$ belongs to $\mathcal{H}$, since f and g are $\mathcal{H}$ -measurable. On $A_{n,k}$ we have $|g| \leq k$, so $\mathbf{1}_{A_{n,k}}g$ is integrable, and the hypothesis gives

$$\mathbb{E}[\mathbf{1}_{A_{n,k}}f] \leq \mathbb{E}[\mathbf{1}_{A_{n,k}}g] \leq k < \infty.$$

Moreover $\mathbb{E}[\mathbf{1}_{A_{n,k}}f] > -\infty$, being non-negative if f is $\overline{\mathbb{R}}_+$ -valued and at least $-\mathbb{E}[f_-] > -\infty$ if f is integrable; hence $\mathbf{1}_{A_{n,k}}f$ is integrable as well. Applying the linearity of expectations (Lemma 2.5.2(ii)),

$$\mathbb{E}[\mathbf{1}_{A_{n,k}}(f - g)] = \mathbb{E}[\mathbf{1}_{A_{n,k}}f] - \mathbb{E}[\mathbf{1}_{A_{n,k}}g] \leq 0.$$

On the other hand $\mathbf{1}_{A_{n,k}}(f - g)$ is a positive random variable with $\mathbf{1}_{A_{n,k}}(f - g) \geq \frac{1}{n}\mathbf{1}_{A_{n,k}}$ pointwise, so Eq. (2.8) gives $\mathbb{E}[\mathbf{1}_{A_{n,k}}(f - g)] \geq \frac{1}{n}\mathbb{P}(A_{n,k})$. Combining the two displays forces $\mathbb{P}(A_{n,k}) = 0$.

Finally, if $f(\omega) > g(\omega)$ and $g(\omega) > -\infty$, then $g(\omega) \in \mathbb{R}$, so $\omega \in A_{n,k}$ for some $n, k \in \mathbb{N}$; that is,

$$\{f > g\} \subseteq \{g = -\infty\} \cup \bigcup_{n,k \in \mathbb{N}} A_{n,k}.$$

The set $\{g = -\infty\}$ is empty if g is $\overline{\mathbb{R}}_+$ -valued and of zero measure if g is integrable, so countable subadditivity (Lemma 2.3.1(v)) yields $\mathbb{P}(f > g) = 0$. $\square$

The following theorem that composes transition probability kernels with probability measures, or other transition probability kernels, via integration, is arguably the most crucial result that upholds the entire theory of causal spaces.

Theorem 2.5.4. *Suppose that K is a transition probability kernel from $(\Omega_1, \mathcal{E}_1)$ into $(\Omega_2, \mathcal{E}_2)$, L a transition probability kernel from $(\Omega_2, \mathcal{E}_2)$ into $(\Omega_3, \mathcal{E}_3)$ and $\mathbb{P}$ a probability measure on $(\Omega_1, \mathcal{E}_1)$.*

(i) *There is a unique probability measure $\mathbb{P} \otimes K$ on $\mathcal{E}_1 \otimes \mathcal{E}_2$ that satisfies*

$$(\mathbb{P} \otimes K)(A \times B) = \int \mathbb{P}(d\omega_1) \mathbf{1}_A(\omega_1) K(\omega_1; B).$$

In particular, the map $\mathcal{E}_2 \rightarrow [0, 1]$ given by $B \mapsto \int \mathbb{P}(d\omega_1) K(\omega_1; B)$ is a probability measure.

(ii) *For any $\mathcal{E}_2$ -measurable function $f : \Omega_2 \rightarrow \overline{\mathbb{R}}_+$, the function $\Omega_1 \rightarrow \overline{\mathbb{R}}_+$ given by $\omega_1 \mapsto \int K(\omega_1; d\omega_2) f(\omega_2)$ is $\mathcal{E}_1$ -measurable.*

(iii) *The function $\Omega_1 \times \mathcal{E}_3 \rightarrow [0, 1]$ given by $(\omega_1, B) \mapsto \int K(\omega_1; d\omega_2) L(\omega_2; B)$ is a transition probability kernel.*

Proof. (i) Define the collection

$$\mathcal{D} = \{C \in \mathcal{E}_1 \otimes \mathcal{E}_2 : \omega_1 \mapsto K(\omega_1; C_{\omega_1}) \text{ is } \mathcal{E}_1\text{-measurable}\},$$

where $C_{\omega_1} \in \mathcal{E}_2$ is the section of C at ω_1 (Eq. (2.1)). Then $\mathcal{D}$ is a λ -system. Indeed, $\omega_1 \mapsto K(\omega_1; (\Omega_1 \times \Omega_2)_{\omega_1}) = K(\omega_1; \Omega_2) = 1$ is a constant function, hence $\mathcal{E}_1$ -measurable, hence $\Omega_1 \times \Omega_2 \in \mathcal{D}$. If $C, D \in \mathcal{D}$ with $C \supseteq D$, then by the additivity of the probability measure $K(\omega_1; \cdot)$, we have that $\omega_1 \mapsto K(\omega_1; (C \setminus D)_{\omega_1}) = K(\omega_1; C_{\omega_1}) - K(\omega_1; D_{\omega_1})$ is $\mathcal{E}_1$ -measurable by Lemma 2.2.2, hence $C \setminus D \in \mathcal{D}$. Finally, if $C_n \in \mathcal{D}$ for all $n \in \mathbb{N}$ with $C_1 \subseteq C_2 \subseteq \dots$, then by Lemma 2.3.1(iii), we have $\omega_1 \mapsto K(\omega_1; (\cup_{n \in \mathbb{N}} C_n)_{\omega_1}) = \lim_{n \rightarrow \infty} K(\omega_1; (C_n)_{\omega_1})$ which is measurable by Theorem 2.2.6(i). Hence $\cup_{n \in \mathbb{N}} C_n \in \mathcal{D}$.

But for rectangles $A \times B$ with $A \in \mathcal{E}_1$ and $B \in \mathcal{E}_2$, $\omega_1 \mapsto K(\omega_1; (A \times B)_{\omega_1}) = \mathbf{1}_A(\omega_1) K(\omega_1; B)$ is $\mathcal{E}_1$ -measurable as a product of two $\mathcal{E}_1$ -measurable functions $\mathbf{1}_A$ and $K(\cdot; B)$ (Lemma 2.2.7). Hence, $A \times B \in \mathcal{D}$. But we know that the collection of rectangles is a π -system that generates $\mathcal{E}_1 \otimes \mathcal{E}_2$, so by the monotone class theorem (Theorem 2.1.5), we have that $\omega_1 \mapsto K(\omega_1; C_{\omega_1})$ is $\mathcal{E}_1$ -measurable for all $C \in \mathcal{E}_1 \otimes \mathcal{E}_2$.

Now define $\mathbb{P} \otimes K : \mathcal{E}_1 \otimes \mathcal{E}_2 \rightarrow [0, 1]$ by $(\mathbb{P} \otimes K)(C) = \int \mathbb{P}(d\omega_1) K(\omega_1; C_{\omega_1})$. Then on rectangles $A \times B$ with $A \in \mathcal{E}_1$ and $B \in \mathcal{E}_2$, we have $(\mathbb{P} \otimes K)(A \times B) = \int \mathbb{P}(d\omega_1) \mathbf{1}_A(\omega_1) K(\omega_1; B)$ as required. It is also clear that $(\mathbb{P} \otimes K)(\Omega_1 \times \Omega_2) = 1$. For σ -additivity, let us take pairwise disjoint $C_1, C_2, \dots \in \mathcal{E}_1 \otimes \mathcal{E}_2$. Then the sections $(C_n)_{\omega_1}$ are also pairwise disjoint and $(\sqcup_{n \in \mathbb{N}} C_n)_{\omega_1} = \sqcup_{n \in \mathbb{N}} (C_n)_{\omega_1}$, so the σ -additivity of $K(\omega_1; \cdot)$ and the monotone convergence theorem (Theorem 2.5.1) give

$$(\mathbb{P} \otimes K)(\sqcup_{n \in \mathbb{N}} C_n) = \int \mathbb{P}(d\omega_1) K(\omega_1; (\sqcup_{n \in \mathbb{N}} C_n)_{\omega_1})$$

$$\begin{aligned}
 &= \int \mathbb{P}(d\omega_1) \sum_{n \in \mathbb{N}} K(\omega_1; (C_n)_{\omega_1}) \\
 &= \sum_{n \in \mathbb{N}} \int \mathbb{P}(d\omega_1) K(\omega_1; (C_n)_{\omega_1}) \\
 &= \sum_{n \in \mathbb{N}} (\mathbb{P} \otimes K)(C_n).
 \end{aligned}$$

Hence $\mathbb{P} \otimes K$ is a probability measure. Its uniqueness is given by Lemma 2.3.2. Finally, the canonical projection $\pi_2 : \Omega_1 \times \Omega_2 \rightarrow \Omega_2$, $(\omega_1, \omega_2) \mapsto \omega_2$, is $(\mathcal{E}_1 \otimes \mathcal{E}_2)/\mathcal{E}_2$ -measurable, since $\pi_2^{-1}(B) = \Omega_1 \times B \in \mathcal{E}_1 \otimes \mathcal{E}_2$ for every $B \in \mathcal{E}_2$. Hence the image measure $(\mathbb{P} \otimes K) \circ \pi_2^{-1}$ is a probability measure on $(\Omega_2, \mathcal{E}_2)$ (see the end of Section 2.3), and putting $A = \Omega_1$, it is given by

$$(\mathbb{P} \otimes K) \circ \pi_2^{-1}(B) = (\mathbb{P} \otimes K)(\Omega_1 \times B) = \int \mathbb{P}(d\omega_1) K(\omega_1; B),$$

which proves the last claim.

- (ii) First, if f is simple and positive, say $f = \sum_{i=1}^n a_i \mathbf{1}_{A_i}$ for some $a_i \in \mathbb{R}_+$ and $A_i \in \mathcal{E}_2$, then the map $\omega_1 \mapsto K(\omega_1; A_i)$ is $\mathcal{E}_1$ -measurable for each $i = 1, \dots, n$, and so as a linear combination of measurable functions, $\omega_1 \mapsto \int K(\omega_1; d\omega_2) f(\omega_2) = \sum_{i=1}^n a_i K(\omega_1; A_i)$ is $\mathcal{E}_1$ -measurable (Lemma 2.2.2).

In general, $f : \Omega_2 \rightarrow \overline{\mathbb{R}}_+$ is $\mathcal{E}_2$ -measurable, so by Theorem 2.2.6(ii), it is the pointwise limit of an increasing sequence of positive simple functions, say $f = \lim_{n \rightarrow \infty} f_n$. Then by the monotone convergence theorem (Theorem 2.5.1),

$$\int K(\omega_1; d\omega_2) f(\omega_2) = \int K(\omega_1; d\omega_2) \lim_{n \rightarrow \infty} f_n(\omega_2) = \lim_{n \rightarrow \infty} \int K(\omega_1; d\omega_2) f_n(\omega_2).$$

Here, each $\int K(\omega_1; d\omega_2) f_n(\omega_2)$ is $\mathcal{E}_1$ -measurable by above, and by the monotonicity property of integrations (Eq. (2.8)), $(\int K(\omega_1; d\omega_2) f_n(\omega_2))_n$ form an increasing sequence. Hence, as a pointwise increasing limit of measurable functions, the map $\omega_1 \mapsto \int K(\omega_1; d\omega_2) f(\omega_2)$ is $\mathcal{E}_1$ -measurable (Theorem 2.2.6(i)).

- (iii) For any $B \in \mathcal{E}_3$, the map $\omega_2 \mapsto L(\omega_2; B)$ is $\mathcal{E}_2$ -measurable, so by (ii), the map $\omega_1 \mapsto \int K(\omega_1; d\omega_2) L(\omega_2; B)$ is $\mathcal{E}_1$ -measurable. On the other hand, for each ω_1 , the map $A \mapsto K(\omega_1; A)$ is a probability measure on $\mathcal{E}_2$, so by (i), the function $B \mapsto \int K(\omega_1; d\omega_2) L(\omega_2; B) : \mathcal{E}_3 \rightarrow [0, 1]$ is a probability measure. Hence the map $(\omega_1, B) \mapsto \int K(\omega_1; d\omega_2) L(\omega_2; B) : \Omega_1 \times \mathcal{E}_3 \rightarrow [0, 1]$ is a transition probability kernel.

□

The resulting transition probability kernel in Theorem 2.5.4(iii) is referred to as the *product* of K and L .

The following measurability property of partial integrals will be used repeatedly when integrating out one argument of a jointly measurable function. Note that it is not a

consequence of Theorem 2.5.4(ii), in which the integrand is allowed to depend on the integration variable only.

Lemma 2.5.5 (Measurability of partial integrals). *Let $(\Omega_1, \mathcal{E}_1)$ and $(\Omega_2, \mathcal{E}_2)$ be measurable spaces, and let $\mathbb{P}$ be a probability measure on $(\Omega_1, \mathcal{E}_1)$. If $f : \Omega_1 \times \Omega_2 \rightarrow \overline{\mathbb{R}}_+$ is $(\mathcal{E}_1 \otimes \mathcal{E}_2)$ -measurable, then the map*

$$\omega_2 \mapsto \int \mathbb{P}(d\omega_1) f((\omega_1, \omega_2))$$

is $\mathcal{E}_2$ -measurable.

Proof. For $C \in \mathcal{E}_1 \otimes \mathcal{E}_2$ and $\omega_2 \in \Omega_2$, recall the sections C_{ω_2} from Eq. (2.1). Define the collection

$$\mathcal{D} = \{C \in \mathcal{E}_1 \otimes \mathcal{E}_2 : \omega_2 \mapsto \mathbb{P}(C_{\omega_2}) \text{ is } \mathcal{E}_2\text{-measurable}\}.$$

Then $\mathcal{D}$ is a λ -system. Indeed, $\omega_2 \mapsto \mathbb{P}((\Omega_1 \times \Omega_2)_{\omega_2}) = \mathbb{P}(\Omega_1) = 1$ is constant, hence $\mathcal{E}_2$ -measurable, so $\Omega_1 \times \Omega_2 \in \mathcal{D}$. If $C, D \in \mathcal{D}$ with $C \supseteq D$, then $(C \setminus D)_{\omega_2} = C_{\omega_2} \setminus D_{\omega_2}$ with $D_{\omega_2} \subseteq C_{\omega_2}$, so by the additivity of $\mathbb{P}$, the map $\omega_2 \mapsto \mathbb{P}((C \setminus D)_{\omega_2}) = \mathbb{P}(C_{\omega_2}) - \mathbb{P}(D_{\omega_2})$ is $\mathcal{E}_2$ -measurable by Lemma 2.2.2, hence $C \setminus D \in \mathcal{D}$. Finally, if $C_n \in \mathcal{D}$ for all $n \in \mathbb{N}$ with $C_1 \subseteq C_2 \subseteq \dots$, then $(\cup_{n \in \mathbb{N}} C_n)_{\omega_2} = \cup_{n \in \mathbb{N}} (C_n)_{\omega_2}$, so by Lemma 2.3.1(iii) we have $\omega_2 \mapsto \mathbb{P}((\cup_{n \in \mathbb{N}} C_n)_{\omega_2}) = \lim_{n \rightarrow \infty} \mathbb{P}((C_n)_{\omega_2})$, which is $\mathcal{E}_2$ -measurable by Theorem 2.2.6(i); hence $\cup_{n \in \mathbb{N}} C_n \in \mathcal{D}$.

But for rectangles $A \times B$ with $A \in \mathcal{E}_1$ and $B \in \mathcal{E}_2$, we have $(A \times B)_{\omega_2} = A$ if $\omega_2 \in B$ and $(A \times B)_{\omega_2} = \emptyset$ otherwise, so

$$\omega_2 \mapsto \mathbb{P}((A \times B)_{\omega_2}) = \mathbb{P}(A) \mathbf{1}_B(\omega_2)$$

is $\mathcal{E}_2$ -measurable, giving $A \times B \in \mathcal{D}$. The rectangles form a π -system that generates $\mathcal{E}_1 \otimes \mathcal{E}_2$, so by the monotone class theorem (Theorem 2.1.5) we obtain $\mathcal{D} = \mathcal{E}_1 \otimes \mathcal{E}_2$. Since $\int \mathbb{P}(d\omega_1) \mathbf{1}_C((\omega_1, \omega_2)) = \mathbb{P}(C_{\omega_2})$, the claim holds for $f = \mathbf{1}_C$ with $C \in \mathcal{E}_1 \otimes \mathcal{E}_2$.

By the linearity of the integral for positive simple functions (Lemma 2.4.3(i)), the claim also holds whenever f is a positive simple function. For general $(\mathcal{E}_1 \otimes \mathcal{E}_2)$ -measurable $f : \Omega_1 \times \Omega_2 \rightarrow \overline{\mathbb{R}}_+$, take positive simple functions $f_n \rightarrow f$ pointwise (Theorem 2.2.6(ii)). The monotone convergence theorem (Theorem 2.5.1) gives, for each $\omega_2 \in \Omega_2$,

$$\int \mathbb{P}(d\omega_1) f((\omega_1, \omega_2)) = \lim_{n \rightarrow \infty} \int \mathbb{P}(d\omega_1) f_n((\omega_1, \omega_2)),$$

and each map $\omega_2 \mapsto \int \mathbb{P}(d\omega_1) f_n((\omega_1, \omega_2))$ is $\mathcal{E}_2$ -measurable by the above. As a pointwise limit of $\mathcal{E}_2$ -measurable functions, the left-hand side is $\mathcal{E}_2$ -measurable by Theorem 2.2.6(i). $\square$

An important consequence of Theorem 2.5.4 is Fubini's theorem, that allows the change of order of integration with respect to a product measure.

Theorem 2.5.6 (Fubini's theorem). *Suppose that $\mathbb{P}$ and $\mathbb{Q}$ are probability measures on $(\Omega_1, \mathcal{E}_1)$ and $(\Omega_2, \mathcal{E}_2)$ respectively. Then for $(\mathcal{E}_1 \otimes \mathcal{E}_2)$ -measurable functions $f : \Omega_1 \times \Omega_2 \rightarrow \overline{\mathbb{R}}_+$ and $(\mathcal{E}_1 \otimes \mathcal{E}_2)$ -measurable and $\mathbb{P} \otimes \mathbb{Q}$ -integrable functions $f : \Omega_1 \times \Omega_2 \rightarrow \overline{\mathbb{R}}$, we have*

$$\int \mathbb{P} \otimes \mathbb{Q}(d\omega) f(\omega) = \int \mathbb{P}(d\omega_1) \int \mathbb{Q}(d\omega_2) f((\omega_1, \omega_2)) = \int \mathbb{Q}(d\omega_2) \int \mathbb{P}(d\omega_1) f((\omega_1, \omega_2)).$$

Proof. We can see that setting $K(\omega_1; \cdot) = \mathbb{Q}(\cdot)$ for all $\omega_1 \in \Omega_1$ in Theorem 2.5.4(i) recovers the product probability measure of Proposition 2.4.4. Now take $C \in \mathcal{E}_1 \otimes \mathcal{E}_2$. Since $\mathbf{1}_{C_{\omega_1}}(\omega_2) = \mathbf{1}_C((\omega_1, \omega_2))$, Theorem 2.5.4(i) reads

$$\int \mathbb{P} \otimes \mathbb{Q}(d\omega) \mathbf{1}_C(\omega) = \int \mathbb{P}(d\omega_1) \mathbb{Q}(C_{\omega_1}) = \int \mathbb{P}(d\omega_1) \int \mathbb{Q}(d\omega_2) \mathbf{1}_C((\omega_1, \omega_2)),$$

so the first claimed identity holds for indicator functions; by the linearity of expectations (Lemma 2.5.2(i)) it holds for non-negative simple functions. For general $\mathcal{E}_1 \otimes \mathcal{E}_2$ -measurable $f : \Omega_1 \times \Omega_2 \rightarrow \overline{\mathbb{R}}_+$, take an increasing sequence $(f_n)_{n \in \mathbb{N}}$ of non-negative simple functions with $f_n \rightarrow f$ pointwise (Theorem 2.2.6(ii)). For each ω_1 , the monotone convergence theorem (Theorem 2.5.1) applied to $f_n((\omega_1, \cdot)) \rightarrow f((\omega_1, \cdot))$ gives

$$\int \mathbb{Q}(d\omega_2) f_n((\omega_1, \omega_2)) \rightarrow \int \mathbb{Q}(d\omega_2) f((\omega_1, \omega_2)),$$

the sequence being increasing by Eq. (2.8); in particular the limit is $\mathcal{E}_1$ -measurable by Theorem 2.2.6(i). A second application of the monotone convergence theorem, to this sequence and to $(f_n)_{n \in \mathbb{N}}$ under $\mathbb{P} \otimes \mathbb{Q}$, gives

$$\begin{aligned} \int \mathbb{P} \otimes \mathbb{Q}(d\omega) f(\omega) &= \lim_{n \rightarrow \infty} \int \mathbb{P}(d\omega_1) \int \mathbb{Q}(d\omega_2) f_n((\omega_1, \omega_2)) \\ &= \int \mathbb{P}(d\omega_1) \int \mathbb{Q}(d\omega_2) f((\omega_1, \omega_2)). \end{aligned}$$

Interchanging the roles of $(\Omega_1, \mathcal{E}_1, \mathbb{P})$ and $(\Omega_2, \mathcal{E}_2, \mathbb{Q})$ gives the second identity.

Finally, let $f : \Omega_1 \times \Omega_2 \rightarrow \overline{\mathbb{R}}$ be $\mathbb{P} \otimes \mathbb{Q}$ -integrable. Applying the above to f_+ and f_- gives

$$\int \mathbb{P}(d\omega_1) \int \mathbb{Q}(d\omega_2) f_{\pm}((\omega_1, \omega_2)) = \int \mathbb{P} \otimes \mathbb{Q}(d\omega) f_{\pm}(\omega) < \infty,$$

so the inner integrals $\int \mathbb{Q}(d\omega_2) f_{\pm}((\omega_1, \omega_2))$ are finite for $\mathbb{P}$ -almost all $\omega_1 \in \Omega_1$; for such ω_1 the inner integral of f is defined and equals their difference. Subtracting the two displays, and using Theorem 2.5.3(i) to discard the $\mathbb{P}$ -null set on which the inner integral is undefined, gives the claim. $\square$

Suppose now that $(\Omega_1, \mathcal{H}_1)$ and $(\Omega_2, \mathcal{H}_2)$ are two measurable spaces, $\mathbb{P}$ a probability measure on $(\Omega_1, \mathcal{H}_1)$ and $f : \Omega_1 \rightarrow \Omega_2$ an $\mathcal{H}_1/\mathcal{H}_2$ -measurable function. Then we saw at the end of Section 2.3 that the image measure $\mathbb{P} \circ f^{-1}$ is a probability measure on $(\Omega_2, \mathcal{H}_2)$. The following lemma is on expectations with respect to the image measure.

Lemma 2.5.7. *Suppose that $g : \Omega_2 \rightarrow \overline{\mathbb{R}}_+$ is an $\mathcal{H}_2$ -measurable function, or that $g : \Omega_2 \rightarrow \overline{\mathbb{R}}$ is $\mathcal{H}_2$ -measurable and $\mathbb{P} \circ f^{-1}$ -integrable. Then*

$$\int \mathbb{P} \circ f^{-1}(d\omega_2) g(\omega_2) = \int \mathbb{P}(d\omega_1) g \circ f(\omega_1).$$

Note that the composition $g \circ f$ is measurable, by the discussion in Section 2.2.

Proof. We will use the fact that a composition of $f : \Omega_1 \rightarrow \Omega_2$ with an indicator function $\mathbf{1}_A : \Omega_2 \rightarrow \{0, 1\}$ for some $A \in \mathcal{H}_2$ is the indicator function $\mathbf{1}_A \circ f = \mathbf{1}_{f^{-1}(A)}$ on Ω_1 with $f^{-1}(A) \in \mathcal{H}_1$.

First, suppose that $g : \Omega_2 \rightarrow \overline{\mathbb{R}}_+$. Recall the simple functions d_n from Eq. (2.4), and write, for simplicity, $d_n \circ g = \sum_{i_n=1}^{m_n} a_{i_n} \mathbf{1}_{A_{i_n}}$ for some $a_{i_n} \in \mathbb{R}_+$ and $A_{i_n} \in \mathcal{H}_2$. Then by the definition of expectations (Eq. (2.6)),

$$\begin{aligned} \int \mathbb{P} \circ f^{-1}(d\omega_2) g(\omega_2) &= \lim_{n \rightarrow \infty} \int \mathbb{P} \circ f^{-1}(d\omega_2) d_n \circ g(\omega_2) \\ &= \lim_{n \rightarrow \infty} \sum_{i_n=1}^{m_n} a_{i_n} \mathbb{P} \circ f^{-1}(A_{i_n}) \\ &= \lim_{n \rightarrow \infty} \sum_{i_n=1}^{m_n} a_{i_n} \int \mathbb{P}(d\omega_1) \mathbf{1}_{f^{-1}(A_{i_n})}(\omega_1) \\ &= \lim_{n \rightarrow \infty} \int \mathbb{P}(d\omega_1) d_n \circ g \circ f(\omega_1) \\ &= \int \mathbb{P}(d\omega_1) g \circ f(\omega_1), \end{aligned}$$

as required.

Now suppose that $g : \Omega_2 \rightarrow \overline{\mathbb{R}}$ is integrable. Since $\max$ commutes with composition, $(g \circ f)_+ = g_+ \circ f$ and $(g \circ f)_- = g_- \circ f$. Applying the $\overline{\mathbb{R}}_+$ case to g_+ and to g_- separately gives

$$\int \mathbb{P}(d\omega_1) (g \circ f)_\pm(\omega_1) = \int \mathbb{P}(d\omega_1) g_\pm \circ f(\omega_1) = \int \mathbb{P} \circ f^{-1}(d\omega_2) g_\pm(\omega_2) < \infty,$$

so $g \circ f$ is integrable with respect to $\mathbb{P}$. Hence, by the definition of the expectation of an $\overline{\mathbb{R}}$ -valued random variable,

$$\begin{aligned} \int \mathbb{P} \circ f^{-1}(d\omega_2) g(\omega_2) &= \int \mathbb{P} \circ f^{-1}(d\omega_2) g_+(\omega_2) - \int \mathbb{P} \circ f^{-1}(d\omega_2) g_-(\omega_2) \\ &= \int \mathbb{P}(d\omega_1) (g \circ f)_+(\omega_1) - \int \mathbb{P}(d\omega_1) (g \circ f)_-(\omega_1) \\ &= \int \mathbb{P}(d\omega_1) g \circ f(\omega_1), \end{aligned}$$

as required. □

Suppose that we have a product measurable space $(\Omega, \mathcal{H}) = \otimes_{t \in \mathbb{T}} (\Omega_t, \mathcal{E}_t)$, and a probability measure $\mathbb{P}$ on $(\Omega, \mathcal{H})$. Then for any $\mathbb{S} \in \mathcal{P}(\mathbb{T})$ and any $\mathcal{E}_{\mathbb{S}}$ -measurable $\mathbb{P} \circ \pi_{\mathbb{S}}^{-1}$ -integrable function $f : \Omega_{\mathbb{S}} \rightarrow \overline{\mathbb{R}}$, we have

$$\int \mathbb{P}(d\omega) f(\omega_{\mathbb{S}}) = \int \mathbb{P}(d\omega) f(\pi_{\mathbb{S}}(\omega)) = \int \mathbb{P} \circ \pi_{\mathbb{S}}^{-1}(d\omega_{\mathbb{S}}) f(\omega_{\mathbb{S}})$$

with a slight abuse of notation.

2.6 Absolute continuity, Radon-Nikodym and Lebesgue decomposition

For two probability measures $\mathbb{P}$ and $\mathbb{Q}$ on a measurable space $(\Omega, \mathcal{H})$, we say that $\mathbb{Q}$ is *absolutely continuous* with respect to $\mathbb{P}$, and write $\mathbb{Q} \ll \mathbb{P}$, if, for all events $A \in \mathcal{H}$ such that $\mathbb{P}(A) = 0$, we have $\mathbb{Q}(A) = 0$. If $\mathcal{G} \subseteq \mathcal{H}$ is a sub- σ -algebra, then we write $\mathbb{Q} \ll_{\mathcal{G}} \mathbb{P}$, and say that $\mathbb{Q}$ is absolutely continuous with respect to $\mathbb{P}$ on $\mathcal{G}$, if, for all $G \in \mathcal{G}$ such that $\mathbb{P}(G) = 0$, we also have $\mathbb{Q}(G) = 0$. If $\mathbb{P} \ll \mathbb{Q}$ and $\mathbb{Q} \ll \mathbb{P}$, then we say $\mathbb{P}$ and $\mathbb{Q}$ are *equivalent*, and write $\mathbb{P} \sim \mathbb{Q}$. If $\mathbb{P} \ll_{\mathcal{G}} \mathbb{Q}$ and $\mathbb{Q} \ll_{\mathcal{G}} \mathbb{P}$, then we write $\mathbb{P} \stackrel{\mathcal{G}}{\sim} \mathbb{Q}$.

If $B \in \mathcal{H}$, then we abuse the notation slightly to write $\mathbb{Q} \ll_{\mathcal{H}|_B} \mathbb{P}$ to mean that, for all $C \in \mathcal{H}|_B$ with $\mathbb{P}(C) = 0$, we have $\mathbb{Q}(C) = 0$. Similarly for $\mathbb{P} \stackrel{\mathcal{H}|_B}{\sim} \mathbb{Q}$.

We present two results based on absolute continuity of measures. First is the *Radon-Nikodym* theorem:

Theorem 2.6.1 (Radon-Nikodym theorem). *If $\mathbb{Q} \ll \mathbb{P}$, then there is a non-negative $\mathcal{H}$ -measurable function $f : \Omega \rightarrow \mathbb{R}_+$ such that, for all $A \in \mathcal{H}$,*

$$\mathbb{Q}(A) = \mathbb{E}[\mathbf{1}_A f].$$

Moreover, f is unique up to $\mathbb{P}$ -null sets, meaning that, for any other $\mathcal{H}$ -measurable function $g : \Omega \rightarrow \mathbb{R}_+$ such that $\mathbb{Q}(A) = \mathbb{E}[\mathbf{1}_A g]$ for all $A \in \mathcal{H}$, we have $f(\omega) = g(\omega)$ for $\mathbb{P}$ -almost all $\omega \in \Omega$.

Proof. We argue existence by a maximisation argument, then uniqueness. Throughout, define the family

$$\mathcal{F} = \{f : \Omega \rightarrow \mathbb{R}_+ : f \text{ is } \mathcal{H}\text{-measurable, and } \mathbb{E}[\mathbf{1}_A f] \leq \mathbb{Q}(A) \text{ for all } A \in \mathcal{H}\},$$

which is non-empty since the zero function lies in $\mathcal{F}$.

Step 1: $\mathcal{F}$ is closed under pairwise maxima. Take $f, g \in \mathcal{F}$, and let $h = \max(f, g)$.

For $A \in \mathcal{H}$, set $A_1 = A \cap \{f \geq g\} \in \mathcal{H}$ and $A_2 = A \setminus A_1$. Then $\mathbf{1}_A h = \mathbf{1}_{A_1} f + \mathbf{1}_{A_2} g$ pointwise, so by the linearity of expectations (Lemma 2.5.2),

$$\mathbb{E}[\mathbf{1}_A h] = \mathbb{E}[\mathbf{1}_{A_1} f] + \mathbb{E}[\mathbf{1}_{A_2} g] \leq \mathbb{Q}(A_1) + \mathbb{Q}(A_2) = \mathbb{Q}(A),$$

hence $h \in \mathcal{F}$.

Step 2: Existence of a maximiser. Let $\alpha = \sup_{g \in \mathcal{F}} \mathbb{E}[g]$; taking $A = \Omega$ in the definition of $\mathcal{F}$ gives $\alpha \leq \mathbb{Q}(\Omega) = 1$. Pick $g_n \in \mathcal{F}$ with $\mathbb{E}[g_n] \rightarrow \alpha$, and set $f_n = \max(g_1, \dots, g_n)$, which lies in $\mathcal{F}$ by induction on Step 1. The sequence $(f_n)_{n \in \mathbb{N}}$ is increasing, with $\mathbb{E}[f_n] \geq \mathbb{E}[g_n]$ by Eq. (2.8). We first check that it is almost surely finite. For $c \in \mathbb{N}$, the pointwise bound $c \mathbf{1}_{\{f_n > c\}} \leq f_n$ and Eq. (2.8) give $\mathbb{P}(f_n > c) \leq \mathbb{E}[f_n]/c \leq 1/c$; as $\{f_n > c\} \uparrow \cup_{n \in \mathbb{N}} \{f_n > c\}$, continuity from below (Lemma 2.3.1(iii)) gives $\mathbb{P}(\cup_{n \in \mathbb{N}} \{f_n > c\}) \leq 1/c$, so the set $B = \cap_{c \in \mathbb{N}} \cup_{n \in \mathbb{N}} \{f_n > c\} \in \mathcal{H}$ on which (f_n) is unbounded satisfies $\mathbb{P}(B) = 0$. Set $f = \lim_{n \rightarrow \infty} f_n \mathbf{1}_{\Omega \setminus B}$, an everywhere-finite positive random variable, $\mathcal{H}$ -measurable by Theorem 2.2.6(i). By the monotone convergence theorem (Theorem 2.5.1) applied to $f_n \mathbf{1}_{\Omega \setminus B} \uparrow f$, together with Theorem 2.5.3(i),

$$\mathbb{E}[f] = \lim_{n \rightarrow \infty} \mathbb{E}[f_n \mathbf{1}_{\Omega \setminus B}] = \lim_{n \rightarrow \infty} \mathbb{E}[f_n] \geq \lim_{n \rightarrow \infty} \mathbb{E}[g_n] = \alpha.$$

For any $A \in \mathcal{H}$, monotone convergence applied to $\mathbf{1}_A f_n \mathbf{1}_{\Omega \setminus B} \uparrow \mathbf{1}_A f$ gives $\mathbb{E}[\mathbf{1}_A f] = \lim_{n \rightarrow \infty} \mathbb{E}[\mathbf{1}_A f_n] \leq \mathbb{Q}(A)$, so $f \in \mathcal{F}$, and hence $\mathbb{E}[f] \leq \alpha$ too. Therefore $\mathbb{E}[f] = \alpha$.

Step 3: The residual measure. Define $\nu : \mathcal{H} \rightarrow [0, 1]$ by

$$\nu(A) = \mathbb{Q}(A) - \mathbb{E}[\mathbf{1}_A f],$$

which is non-negative since $f \in \mathcal{F}$. For disjoint $A_1, A_2, \dots \in \mathcal{H}$, the partial sums $\sum_{n=1}^m \mathbf{1}_{A_n} f$ form an increasing sequence of positive random variables converging pointwise to $\mathbf{1}_{\cup_{n \in \mathbb{N}} A_n} f$, so we can apply the monotone convergence theorem (Theorem 2.5.1) and the linearity of expectations (Lemma 2.5.2) to obtain

$$\mathbb{E}[\mathbf{1}_{\cup_{n \in \mathbb{N}} A_n} f] = \lim_{m \rightarrow \infty} \sum_{n=1}^m \mathbb{E}[\mathbf{1}_{A_n} f] = \sum_{n \in \mathbb{N}} \mathbb{E}[\mathbf{1}_{A_n} f].$$

Combined with the σ -additivity of $\mathbb{Q}$, this gives σ -additivity of ν . We claim $\nu(\Omega) = 0$, which by monotonicity of ν yields $\nu \equiv 0$, i.e. $\mathbb{Q}(A) = \mathbb{E}[\mathbf{1}_A f]$ for all $A \in \mathcal{H}$ as required.

Step 4: $\nu(\Omega) = 0$ by contradiction. Suppose for contradiction $\nu(\Omega) > 0$. Pick $n \in \mathbb{N}$ with $\nu(\Omega) > 1/n$, and define

$$\mu(A) = \nu(A) - \frac{1}{n} \mathbb{P}(A), \quad A \in \mathcal{H},$$

which takes values in $[-1/n, 1]$ with $\mu(\Omega) > 0$. We construct $E \in \mathcal{H}$ with $\mathbb{P}(E) > 0$ such that $\mu(B) \geq 0$ for every $B \in \mathcal{H}$ with $B \subseteq E$ —a *positive set* for μ .

Set $E_0 = \Omega$, and proceed recursively: given E_k , let

$$\delta_k = \inf\{\mu(B) : B \in \mathcal{H}, B \subseteq E_k\} \in [-1/n, 0].$$

If $\delta_k = 0$, terminate. Otherwise pick $B_k \subseteq E_k$ in $\mathcal{H}$ with $\mu(B_k) < \delta_k/2 < 0$, and set $E_{k+1} = E_k \setminus B_k$. Since B_k and E_{k+1} are disjoint with union E_k , additivity

of μ gives $\mu(E_{k+1}) = \mu(E_k) - \mu(B_k) > \mu(E_k) + |\delta_k|/2$. The sequence $(\mu(E_k))_k$ is therefore strictly increasing while the recursion continues, and bounded above by $\nu(\Omega) \leq 1$, so it converges, forcing $\delta_k \rightarrow 0$.

If the recursion terminates at some k , set $E = E_k$; then $\mu(B) \geq \delta_k = 0$ for every $B \subseteq E$ in $\mathcal{H}$, and $\mu(E) \geq \mu(E_0) = \mu(\Omega) > 0$. Otherwise set $E = \bigcap_{k \in \mathbb{N}} E_k$; since ν and $\mathbb{P}$ are both finite measures, the continuity-from-above property of Lemma 2.3.1(iv) gives $\mu(E) = \lim_{k \rightarrow \infty} \mu(E_k) \geq \mu(\Omega) > 0$ (Lemma 2.3.1 is stated for probability measures, but the proof goes through verbatim without requiring normalisation to 1), and for any $B \subseteq E$ in $\mathcal{H}$, $B \subseteq E_k$ gives $\mu(B) \geq \delta_k \rightarrow 0$, hence $\mu(B) \geq 0$. In either case, E is a positive set for μ with $\mu(E) > 0$.

We claim $\mathbb{P}(E) > 0$. If instead $\mathbb{P}(E) = 0$, then $\mathbb{Q}(E) = 0$ by $\mathbb{Q} \ll \mathbb{P}$, so $\nu(E) \leq \mathbb{Q}(E) = 0$, giving $\mu(E) = \nu(E) - \frac{1}{n}\mathbb{P}(E) = 0$, contradicting $\mu(E) > 0$.

Now set $g = f + \frac{1}{n}\mathbf{1}_E$, an everywhere-finite positive random variable. For any $A \in \mathcal{H}$, Lemma 2.5.2 gives

$$\begin{aligned} \mathbb{E}[\mathbf{1}_A g] &= \mathbb{E}[\mathbf{1}_A f] + \frac{1}{n}\mathbb{P}(A \cap E) \\ &= \mathbb{Q}(A) - \nu(A) + \frac{1}{n}\mathbb{P}(A \cap E) \\ &\leq \mathbb{Q}(A) - \nu(A \cap E) + \frac{1}{n}\mathbb{P}(A \cap E) \\ &= \mathbb{Q}(A) - \mu(A \cap E) \\ &\leq \mathbb{Q}(A), \end{aligned}$$

using $\nu(A) \geq \nu(A \cap E)$ (since ν is non-negative) and $\mu(A \cap E) \geq 0$ (since $A \cap E \subseteq E$ and E is a positive set for μ). Hence $g \in \mathcal{F}$, but

$$\mathbb{E}[g] = \mathbb{E}[f] + \frac{1}{n}\mathbb{P}(E) = \alpha + \frac{1}{n}\mathbb{P}(E) > \alpha,$$

contradicting the maximality of α . This contradiction shows $\nu(\Omega) = 0$, completing the existence argument.

Step 5: Uniqueness. If $f, g : \Omega \rightarrow \mathbb{R}_+$ both satisfy $\mathbb{Q}(A) = \mathbb{E}[\mathbf{1}_A f] = \mathbb{E}[\mathbf{1}_A g]$ for all $A \in \mathcal{H}$, then Theorem 2.5.3(ii) gives $f = g$ $\mathbb{P}$ -almost surely.

□

We write $\frac{d\mathbb{Q}}{d\mathbb{P}}$ for f and call it the *Radon-Nikodym derivative*. Next, we present the *Lebesgue decomposition* of two probability measures.

Theorem 2.6.2 (Lebesgue decomposition). *Let $\mathbb{P}$ and $\mathbb{Q}$ be two probability measures on a measurable space $(\Omega, \mathcal{H})$. There exist measurable subsets $\Lambda^1, \Lambda^2, \Lambda^3 \subseteq \Omega$ such that*

- (i) $\Omega = \Lambda^1 \sqcup \Lambda^2 \sqcup \Lambda^3$;
- (ii) $\mathbb{P}(\Lambda^1) = \mathbb{Q}(\Lambda^3) = 0$;
- (iii) $\mathbb{P}|_{\Lambda^2} \sim \mathbb{Q}|_{\Lambda^2}$.

Moreover, such a partition is unique up to $(\mathbb{P} + \mathbb{Q})$ -null sets: if $\tilde{\Lambda}^1, \tilde{\Lambda}^2, \tilde{\Lambda}^3 \subseteq \Omega$ also satisfy the above properties, then

$$\frac{1}{2}(\mathbb{P} + \mathbb{Q})(\Lambda^1 \triangle \tilde{\Lambda}^1) = \frac{1}{2}(\mathbb{P} + \mathbb{Q})(\Lambda^2 \triangle \tilde{\Lambda}^2) = \frac{1}{2}(\mathbb{P} + \mathbb{Q})(\Lambda^3 \triangle \tilde{\Lambda}^3) = 0.$$

Finally, Λ^2 is the maximal (up to $\mathbb{P}$ -null sets) set on which $\mathbb{P}$ and $\mathbb{Q}$ are mutually absolutely continuous.

Proof. We first show existence. Let $\mu = \frac{1}{2}(\mathbb{P} + \mathbb{Q})$, which is a probability measure on $(\Omega, \mathcal{H})$ since $\mathbb{P}$ and $\mathbb{Q}$ are. For any $A \in \mathcal{H}$ with $\mu(A) = 0$, we have $\mathbb{P}(A) + \mathbb{Q}(A) = 2\mu(A) = 0$, forcing $\mathbb{P}(A) = \mathbb{Q}(A) = 0$; in particular, $\mathbb{P} \ll \mu$ and $\mathbb{Q} \ll \mu$. By the Radon-Nikodym theorem (Theorem 2.6.1), there exist $\mathcal{H}$ -measurable functions $f, g : \Omega \rightarrow \mathbb{R}_+$ such that, for all $A \in \mathcal{H}$,

$$\mathbb{P}(A) = \int \mu(d\omega) \mathbf{1}_A(\omega) f(\omega), \quad \mathbb{Q}(A) = \int \mu(d\omega) \mathbf{1}_A(\omega) g(\omega).$$

Define the $\mathcal{H}$ -measurable sets

$$\begin{aligned} \Lambda^1 &= \{\omega \in \Omega : f(\omega) = 0, g(\omega) > 0\}, \\ \Lambda^2 &= \{\omega \in \Omega : f(\omega) > 0, g(\omega) > 0\}, \\ \Lambda^3 &= \{\omega \in \Omega : g(\omega) = 0\}. \end{aligned}$$

These are pairwise disjoint with union Ω , establishing (i). For (ii),

$$\mathbb{P}(\Lambda^1) = \int \mu(d\omega) \mathbf{1}_{\Lambda^1}(\omega) f(\omega) = 0$$

since $\mathbf{1}_{\Lambda^1} f$ vanishes pointwise on Ω ; the equality $\mathbb{Q}(\Lambda^3) = 0$ follows by the symmetric argument.

For (iii), we show that, for any $A \in \mathcal{H}$,

$$\mathbb{P}(A \cap \Lambda^2) = 0 \iff \mu(A \cap \Lambda^2) = 0 \iff \mathbb{Q}(A \cap \Lambda^2) = 0,$$

which gives the mutual absolute continuity of $\mathbb{P}$ and $\mathbb{Q}$ restricted on Λ^2 . Fix $A \in \mathcal{H}$ and write $B = A \cap \Lambda^2$.

If $\mu(B) = 0$, then $\mathbf{1}_B f$ vanishes μ -almost everywhere, so by Theorem 2.5.3(i), $\mathbb{P}(B) = \int \mathbf{1}_B(\omega) f(\omega) \mu(d\omega) = 0$. Conversely, suppose $\mathbb{P}(B) = 0$. For each $n \in \mathbb{N}$, let $B_n = B \cap \{f \geq \frac{1}{n}\} \in \mathcal{H}$; since $f > 0$ on $B \subseteq \Lambda^2$, we have $B = \bigcup_{n \in \mathbb{N}} B_n$. The pointwise bound $\mathbf{1}_{B_n} f \geq \frac{1}{n} \mathbf{1}_{B_n}$ on Ω and monotonicity (Eq. (2.8)) applied under μ give

$$0 = \mathbb{P}(B) \geq \mathbb{P}(B_n) = \int \mu(d\omega) \mathbf{1}_{B_n}(\omega) f(\omega) \geq \frac{1}{n} \mu(B_n),$$

so $\mu(B_n) = 0$ for each n , and $\mu(B) = \mu(\bigcup_{n \in \mathbb{N}} B_n) \leq \sum_{n \in \mathbb{N}} \mu(B_n) = 0$, by countable subadditivity (Lemma 2.3.1(v)). The chain involving $\mathbb{Q}$ follows by the same argument with g in place of f , using that $g > 0$ on Λ^2 .

We now show uniqueness. Suppose $\tilde{\Lambda}^1, \tilde{\Lambda}^2, \tilde{\Lambda}^3$ is another partition satisfying (i)–(iii). Note that $\mu(B) = 0$ if and only if both $\mathbb{P}(B) = 0$ and $\mathbb{Q}(B) = 0$, since $\mathbb{P}, \mathbb{Q} \geq 0$.

For Λ^1 , monotonicity gives $\mathbb{P}(\Lambda^1 \triangle \tilde{\Lambda}^1) \leq \mathbb{P}(\Lambda^1) + \mathbb{P}(\tilde{\Lambda}^1) = 0$. For the $\mathbb{Q}$ -measure, decompose

$$\Lambda^1 \setminus \tilde{\Lambda}^1 = (\Lambda^1 \cap \tilde{\Lambda}^2) \sqcup (\Lambda^1 \cap \tilde{\Lambda}^3).$$

Then $\mathbb{P}(\Lambda^1 \cap \tilde{\Lambda}^2) \leq \mathbb{P}(\Lambda^1) = 0$, and the mutual absolute continuity of $\mathbb{P}$ and $\mathbb{Q}$ on $\tilde{\Lambda}^2$ gives $\mathbb{Q}(\Lambda^1 \cap \tilde{\Lambda}^2) = 0$. Also $\mathbb{Q}(\Lambda^1 \cap \tilde{\Lambda}^3) \leq \mathbb{Q}(\tilde{\Lambda}^3) = 0$. Hence $\mathbb{Q}(\Lambda^1 \setminus \tilde{\Lambda}^1) = 0$, and by interchanging the roles of Λ^1 and $\tilde{\Lambda}^1$ in the argument, $\mathbb{Q}(\tilde{\Lambda}^1 \setminus \Lambda^1) = 0$, so $\mathbb{Q}(\Lambda^1 \triangle \tilde{\Lambda}^1) = 0$. Therefore $\mu(\Lambda^1 \triangle \tilde{\Lambda}^1) = 0$.

The result for Λ^3 follows by interchanging the roles of $\mathbb{P}$ and $\mathbb{Q}$ (and of Λ^1 and Λ^3) in the argument above.

For Λ^2 , decompose

$$\Lambda^2 \setminus \tilde{\Lambda}^2 = (\Lambda^2 \cap \tilde{\Lambda}^1) \sqcup (\Lambda^2 \cap \tilde{\Lambda}^3).$$

Now $\mathbb{P}(\Lambda^2 \cap \tilde{\Lambda}^1) \leq \mathbb{P}(\tilde{\Lambda}^1) = 0$, and mutual absolute continuity on Λ^2 gives $\mathbb{Q}(\Lambda^2 \cap \tilde{\Lambda}^1) = 0$; similarly, $\mathbb{Q}(\Lambda^2 \cap \tilde{\Lambda}^3) \leq \mathbb{Q}(\tilde{\Lambda}^3) = 0$, and mutual absolute continuity on Λ^2 gives $\mathbb{P}(\Lambda^2 \cap \tilde{\Lambda}^3) = 0$. Hence $\mu(\Lambda^2 \setminus \tilde{\Lambda}^2) = 0$, and by symmetry $\mu(\tilde{\Lambda}^2 \setminus \Lambda^2) = 0$, yielding $\mu(\Lambda^2 \triangle \tilde{\Lambda}^2) = 0$.

Finally, for maximality, suppose that $\Lambda \in \mathcal{H}$ with $\mathbb{P} \stackrel{\mathcal{H}|\Lambda}{\sim} \mathbb{Q}$, then $\Lambda \setminus \Lambda^2 = (\Lambda \cap \Lambda^1) \sqcup (\Lambda \cap \Lambda^3)$. Here, the first has $\mathbb{P}$ -measure zero by (ii), and the second has $\mathbb{Q}$ -measure zero by (ii), and hence zero $\mathbb{P}$ -measure by $\mathbb{P} \ll \mathbb{Q}$ on $\mathcal{H}|\Lambda$. $\square$

2.7 Conditioning

The theory of conditioning plays a crucial role in the study of causality. Let $G \in \mathcal{H}$ be an event, and $\mathcal{G} \subseteq \mathcal{H}$ a sub- σ -algebra.

For an event $A \in \mathcal{H}$, the *conditional probability* of A given G is defined as

$$\mathbb{P}_G(A) = \begin{cases} \frac{\mathbb{P}(G \cap A)}{\mathbb{P}(G)} & \text{if } \mathbb{P}(G) > 0 \\ \text{undefined} & \text{otherwise.} \end{cases}$$

Note that we use a somewhat uncommon notation $\mathbb{P}_G(\cdot)$ rather than the more common $\mathbb{P}(\cdot | G)$. We will discuss this after we introduce conditioning given σ -algebras $\mathcal{G}$.

We define the *conditional expectation* of a random variable $f : \Omega \rightarrow \mathbb{R}$ given G as

$$\mathbb{E}_G[f] = \begin{cases} \frac{\mathbb{E}[\mathbf{1}_G f]}{\mathbb{P}(G)} & \text{if } \mathbb{P}(G) > 0 \\ \text{undefined} & \text{otherwise.} \end{cases}$$

We say that a finite number of events $A_1, \dots, A_n \in \mathcal{H}$ are *conditionally independent* given G if $\mathbb{P}(G) > 0$ and, for any non-empty subset $I \subseteq \{1, \dots, n\}$,

$$\mathbb{P}_G(\cap_{i \in I} A_i) = \prod_{i \in I} \mathbb{P}_G(A_i).$$

If $\mathbb{P}(G) = 0$, then the conditional independence given G of any events cannot be determined. We say that a finite number of sub- σ -algebras $\mathcal{H}_1, \dots, \mathcal{H}_n \subseteq \mathcal{H}$ are conditionally independent given G if all tuples of events $A_1 \in \mathcal{H}_1, \dots, A_n \in \mathcal{H}_n$ are conditionally independent given G . Finally, an arbitrary set $\{\mathcal{H}_t\}_{t \in \mathbb{T}}$ of sub- σ -algebras of $\mathcal{H}$, where $\mathbb{T}$ can be uncountably infinite, are said to be conditionally independent given G if every finite subset of them are conditionally independent given G .

The more involved case is when we condition on a sub- σ -algebra $\mathcal{G}$. Let us first consider a positive random variable $f : \Omega \rightarrow \mathbb{R}_+$. The *conditional expectation* of f given $\mathcal{G}$ is defined as any $\mathcal{G}$ -measurable random variable $\gamma \mapsto \mathbb{E}_{\mathcal{G}}[\gamma, f] : \Omega \rightarrow \mathbb{R}_+$ that satisfies, for any $G \in \mathcal{G}$,

$$\mathbb{E}[\mathbf{1}_G f] = \mathbb{E}[\mathbf{1}_G(\cdot) \mathbb{E}_{\mathcal{G}}[\cdot, f]]. \quad (2.10)$$

In particular, the *conditional probability* of an event $A \in \mathcal{H}$ given $\mathcal{G}$ is defined as any $\mathcal{G}$ -measurable random variable $\gamma \mapsto \mathbb{P}_{\mathcal{G}}(\gamma, A) : \Omega \rightarrow [0, 1]$ such that, for all $G \in \mathcal{G}$,

$$\mathbb{P}(G \cap A) = \mathbb{E}[\mathbf{1}_G(\cdot) \mathbb{P}_{\mathcal{G}}(\cdot, A)],$$

obtained by letting $f = \mathbf{1}_A$ above. The conditional expectation of an arbitrary random variable f given $\mathcal{G}$ is defined as the $\mathcal{G}$ -measurable random variable

$$\mathbb{E}_{\mathcal{G}}[\cdot, f] = \begin{cases} \mathbb{E}_{\mathcal{G}}[\cdot, f_+] - \mathbb{E}_{\mathcal{G}}[\cdot, f_-] & \text{if } f \text{ is integrable (see Section 2.5)} \\ \text{undefined} & \text{otherwise,} \end{cases}$$

where, as in Section 2.5, we used the positive random variables $f_+ = \max\{f, 0\}$ and $f_- = -\min\{f, 0\}$.

Note that $\mathbb{E}_G[f], \mathbb{P}_G(A)$ and $\mathbb{E}_{\mathcal{G}}[\cdot, f], \mathbb{P}_{\mathcal{G}}(\cdot, A)$ are different types of objects—the former are numbers, whereas the latter are $\mathcal{G}$ -measurable random variables.

Since the definition is not constructive, we must show that the conditional expectation given $\mathcal{G}$ exists, and is (almost surely) unique.

Theorem 2.7.1 (Unique existence of conditional expectation). *The conditional expectation $\mathbb{E}_{\mathcal{G}}[\cdot, f]$ exists if f is non-negative or integrable, and is unique $\mathbb{P}$ -almost surely.*

Proof. We argue existence in the bounded non-negative case, then in the general non-negative case, then in the integrable case, and finally treat uniqueness.

Step 1: Existence for bounded non-negative f . Suppose first that $0 \leq f \leq c$ for some $c \geq 0$, and define the set function

$$\nu(G) = \mathbb{E}[\mathbf{1}_G f], \quad G \in \mathcal{G}.$$

For disjoint $G_1, G_2, \dots \in \mathcal{G}$, the partial sums $\sum_{n=1}^m \mathbf{1}_{G_n} f$ form an increasing sequence of positive random variables converging pointwise to $\mathbf{1}_{\sqcup_{n \in \mathbb{N}} G_n} f$, so the monotone convergence theorem (Theorem 2.5.1) and the linearity of expectations (Lemma 2.5.2) give

$$\nu(\sqcup_{n \in \mathbb{N}} G_n) = \mathbb{E} \left[\lim_{m \rightarrow \infty} \sum_{n=1}^m \mathbf{1}_{G_n} f \right] = \lim_{m \rightarrow \infty} \sum_{n=1}^m \mathbb{E}[\mathbf{1}_{G_n} f] = \sum_{n \in \mathbb{N}} \nu(G_n).$$

Hence, ν is a σ -additive set function, with $\nu(\emptyset) = 0$ and $\nu(\Omega) \leq c$. Moreover, $\nu \ll_{\mathcal{G}} \mathbb{P}$: if $G \in \mathcal{G}$ satisfies $\mathbb{P}(G) = 0$, then $\mathbb{E}[\mathbf{1}_G f] \leq c\mathbb{P}(G) = 0$, so $\nu(G) = 0$.

If $\nu(\Omega) = 0$, then $\nu(G) \leq \nu(\Omega) = 0$ for every $G \in \mathcal{G}$, meaning that $\mathbb{E}[\mathbf{1}_G f] = 0 = \mathbb{E}[\mathbf{1}_G \cdot 0]$, so $g = 0$ is a version of $\mathbb{E}_{\mathcal{G}}[\cdot, f]$. Otherwise, set $\alpha = \nu(\Omega) > 0$ and define the probability measure $\mathbb{Q} = \frac{\nu}{\alpha}$ on $(\Omega, \mathcal{G})$, which still satisfies $\mathbb{Q} \ll_{\mathcal{G}} \mathbb{P}$. Applying the Radon-Nikodym theorem (Theorem 2.6.1) to $\mathbb{P}$ and $\mathbb{Q}$ on the measurable space $(\Omega, \mathcal{G})$ gives a non-negative $\mathcal{G}$ -measurable function $h : \Omega \rightarrow \mathbb{R}_+$ with $\mathbb{Q}(G) = \mathbb{E}[\mathbf{1}_G h]$ for all $G \in \mathcal{G}$. Setting $g = \alpha h$ yields a non-negative $\mathcal{G}$ -measurable function with

$$\mathbb{E}[\mathbf{1}_G f] = \nu(G) = \alpha \mathbb{Q}(G) = \mathbb{E}[\mathbf{1}_G g] \quad \text{for all } G \in \mathcal{G}.$$

Step 2: Existence for non-negative f . Now let $f : \Omega \rightarrow \overline{\mathbb{R}}_+$ be arbitrary. For each $n \in \mathbb{N}$, $f_n = \min\{f, n\}$ is bounded, so by Step 1 there exists a non-negative $\mathcal{G}$ -measurable $g_n : \Omega \rightarrow \mathbb{R}_+$ with $\mathbb{E}[\mathbf{1}_G f_n] = \mathbb{E}[\mathbf{1}_G g_n]$ for all $G \in \mathcal{G}$. For each n , and any $G \in \mathcal{G}$,

$$\mathbb{E}[\mathbf{1}_G (g_{n+1} - g_n)] = \mathbb{E}[\mathbf{1}_G (f_{n+1} - f_n)] \geq 0.$$

Taking $G = \{g_{n+1} < g_n\} \in \mathcal{G}$, where the integrand is ≤ 0 , forces $\mathbb{P}(\{g_{n+1} < g_n\}) = 0$, so $g_n \leq g_{n+1}$ almost surely. Modifying each g_n on a $\mathbb{P}$ -null set in $\mathcal{G}$ if necessary, we may assume the inequality holds pointwise. Define

$$g = \lim_{n \rightarrow \infty} g_n,$$

which is $\mathcal{G}$ -measurable and takes values in $[0, \infty]$ by Theorem 2.2.6(i). For each $G \in \mathcal{G}$, the monotone convergence theorem (Theorem 2.5.1) applied to $\mathbf{1}_G g_n \uparrow \mathbf{1}_G g$ and $\mathbf{1}_G f_n \uparrow \mathbf{1}_G f$ gives

$$\mathbb{E}[\mathbf{1}_G g] = \lim_{n \rightarrow \infty} \mathbb{E}[\mathbf{1}_G g_n] = \lim_{n \rightarrow \infty} \mathbb{E}[\mathbf{1}_G f_n] = \mathbb{E}[\mathbf{1}_G f].$$

Hence g is a version of $\mathbb{E}_{\mathcal{G}}[\cdot, f]$.

Step 3: Existence for integrable f . Let $f : \Omega \rightarrow \overline{\mathbb{R}}$ be integrable, and decompose $f = f_+ - f_-$. By Step 2 there exist non-negative $\mathcal{G}$ -measurable $g_{\pm}$ with $\mathbb{E}[\mathbf{1}_G f_{\pm}] = \mathbb{E}[\mathbf{1}_G g_{\pm}]$ for all $G \in \mathcal{G}$. Setting $G = \Omega$ gives $\mathbb{E}[g_{\pm}] = \mathbb{E}[f_{\pm}] < \infty$, so $g_{\pm} < \infty$ almost surely; redefining each $g_{\pm}$ to be 0 on the $\mathcal{G}$ -measurable null set $\{g_{\pm} = \infty\}$ preserves $\mathbb{E}[\mathbf{1}_G g_{\pm}]$ by Theorem 2.5.3(i), so we may assume $g_{\pm} : \Omega \rightarrow \mathbb{R}_+$ are everywhere finite. Set $g = g_+ - g_-$, which is $\mathcal{G}$ -measurable, and apply linearity of expectations (Lemma 2.5.2):

$$\mathbb{E}[\mathbf{1}_G f] = \mathbb{E}[\mathbf{1}_G f_+] - \mathbb{E}[\mathbf{1}_G f_-] = \mathbb{E}[\mathbf{1}_G g_+] - \mathbb{E}[\mathbf{1}_G g_-] = \mathbb{E}[\mathbf{1}_G g]$$

for every $G \in \mathcal{G}$, matching the definition of $\mathbb{E}_{\mathcal{G}}[\cdot, f]$ in the integrable case.

Step 4: Uniqueness. Let g, g' be two versions of $\mathbb{E}_{\mathcal{G}}[\cdot, f]$, so that both are $\mathcal{G}$ -measurable and $\mathbb{E}[\mathbf{1}_G g] = \mathbb{E}[\mathbf{1}_G f] = \mathbb{E}[\mathbf{1}_G g']$ for all $G \in \mathcal{G}$. Theorem 2.5.3(ii) immediately gives that $g = g'$ almost surely.

□

Since the uniqueness is only guaranteed $\mathbb{P}$ -almost surely, we call each possible realisation of the conditional expectation $\mathbb{E}_{\mathcal{G}}[\cdot, f]$ a *version*.

We use the somewhat unusual notation $\mathbb{P}_{\mathcal{G}}(\cdot, A)$ and $\mathbb{E}_{\mathcal{G}}[\cdot, f]$ for the conditional objects, rather than the more common $\mathbb{P}(A \mid \mathcal{G})$ and $\mathbb{E}[f \mid \mathcal{G}]$. This is because very often in this monograph, the conditional probability will be compared against a transition probability kernel $K_{\mathbb{S}}(\cdot; A)$, and we want the notation of the conditional probability to parallel that of transition probability kernels. We note that this notation is also sometimes in use in standard probability theory textbooks ([Çinlar, 2011](#)).

Here are some properties of conditional expectations that are similar to ordinary expectations, namely, monotonicity, linearity and monotone convergence.

Lemma 2.7.2. *Let $\mathcal{G}$ be a sub- σ -algebra of $\mathcal{H}$. All functions in this lemma are $\mathcal{H}$ -measurable, and f, g are either both $\mathbb{R}_+$ -valued or both $\mathbb{R}$ -valued and integrable, so that all the expectations exist.*

- (i) *If $f \leq g$, then $\mathbb{E}_{\mathcal{G}}[\cdot, f] \leq \mathbb{E}_{\mathcal{G}}[\cdot, g]$ almost surely.*
- (ii) *If f, g are $\mathbb{R}_+$ -valued and $a, b \in \mathbb{R}_+$, or if f, g are integrable and $a, b \in \mathbb{R}$ with $af + bg$ well-defined, then $\mathbb{E}_{\mathcal{G}}[\cdot, af + bg] = a\mathbb{E}_{\mathcal{G}}[\cdot, f] + b\mathbb{E}_{\mathcal{G}}[\cdot, g]$ almost surely.*
- (iii) *If $(f_n)_{n \in \mathbb{N}}$ is an increasing sequence of $\mathbb{R}_+$ -valued functions with limit f , then $\lim_{n \rightarrow \infty} \mathbb{E}_{\mathcal{G}}[\cdot, f_n]$ exists and equals $\mathbb{E}_{\mathcal{G}}[\cdot, f]$ almost surely.*

Proof. (i) Fix $G \in \mathcal{G}$. Since $\mathbf{1}_G \geq 0$, the hypothesis $f \leq g$ gives $\mathbf{1}_G f \leq \mathbf{1}_G g$ pointwise, whence $(\mathbf{1}_G f)_+ \leq (\mathbf{1}_G g)_+$ and $(\mathbf{1}_G f)_- \geq (\mathbf{1}_G g)_-$; two applications of Eq. (2.8) then give $\mathbb{E}[\mathbf{1}_G f] \leq \mathbb{E}[\mathbf{1}_G g]$. But by the definition of conditional expectations,

$$\mathbb{E}[\mathbf{1}_G \mathbb{E}_{\mathcal{G}}[\cdot, f]] = \mathbb{E}[\mathbf{1}_G f] \leq \mathbb{E}[\mathbf{1}_G g] = \mathbb{E}[\mathbf{1}_G \mathbb{E}_{\mathcal{G}}[\cdot, g]] \quad \text{for every } G \in \mathcal{G}.$$

Now apply Theorem 2.5.3(iii).

- (ii) For each $G \in \mathcal{G}$, we can apply Lemma 2.5.2 to see that

$$\begin{aligned} \mathbb{E}[\mathbf{1}_G(af + bg)] &= a\mathbb{E}[\mathbf{1}_G f] + b\mathbb{E}[\mathbf{1}_G g] \\ &= a\mathbb{E}[\mathbf{1}_G \mathbb{E}_{\mathcal{G}}[\cdot, f]] + b\mathbb{E}[\mathbf{1}_G \mathbb{E}_{\mathcal{G}}[\cdot, g]] \\ &= \mathbb{E}[\mathbf{1}_G(a\mathbb{E}_{\mathcal{G}}[\cdot, f] + b\mathbb{E}_{\mathcal{G}}[\cdot, g])], \end{aligned}$$

so by applying Theorem 2.5.3(ii), we have $\mathbb{E}_{\mathcal{G}}[\cdot, af + bg] = a\mathbb{E}_{\mathcal{G}}[\cdot, f] + b\mathbb{E}_{\mathcal{G}}[\cdot, g]$ almost surely.

- (iii) First note that, for each $n \in \mathbb{N}$, we have $\mathbb{E}_{\mathcal{G}}[\cdot, f_n] \leq \mathbb{E}_{\mathcal{G}}[\cdot, f_{n+1}]$ almost surely by part (i), which means that, by σ -additivity, there is an event in $\mathcal{G}$ of measure 1 on which $\mathbb{E}_{\mathcal{G}}[\cdot, f_n] \leq \mathbb{E}_{\mathcal{G}}[\cdot, f_{n+1}]$ for all $n \in \mathbb{N}$ simultaneously. For the rest of the proof, we stay on this event. For any $G \in \mathcal{G}$, we can apply the monotone convergence theorem (Theorem 2.5.1) to see that

$$\mathbb{E}[\mathbf{1}_G \mathbb{E}_{\mathcal{G}}[\cdot, f]] = \mathbb{E}[\mathbf{1}_G f]$$

$$\begin{aligned}
 &= \lim_{n \rightarrow \infty} \mathbb{E}[\mathbf{1}_G f_n] \\
 &= \lim_{n \rightarrow \infty} \mathbb{E}[\mathbf{1}_G \mathbb{E}_{\mathcal{G}}[\cdot, f_n]] \\
 &= \mathbb{E}[\mathbf{1}_G \lim_{n \rightarrow \infty} \mathbb{E}_{\mathcal{G}}[\cdot, f_n]],
 \end{aligned}$$

so by applying Theorem 2.5.3(ii), we have $\mathbb{E}_{\mathcal{G}}[\cdot, f] = \lim_{n \rightarrow \infty} \mathbb{E}_{\mathcal{G}}[\cdot, f_n]$ almost surely. $\square$

The following are equivalent definitions of conditional expectations.

Proposition 2.7.3. *Let $f : \Omega \rightarrow \overline{\mathbb{R}}_+$ be a positive random variable, or let $f : \Omega \rightarrow \overline{\mathbb{R}}$ be integrable.*

(i) *For any $\mathcal{G}$ -measurable random variable $h : \Omega \rightarrow \overline{\mathbb{R}}_+$ such that hf is non-negative or integrable, we have $\mathbb{E}[hf] = \mathbb{E}[h\mathbb{E}_{\mathcal{G}}[\cdot, f]]$.*

Conversely, if a $\mathcal{G}$ -measurable random variable g satisfies $\mathbb{E}[hf] = \mathbb{E}[hg]$ for every such h , then g is a version of $\mathbb{E}_{\mathcal{G}}[\cdot, f]$.

(ii) *Let $\mathcal{C}$ be a π -system with $\Omega \in \mathcal{C}$ and $\mathcal{G} = \sigma(\mathcal{C})$, and let $f : \Omega \rightarrow \overline{\mathbb{R}}$ be integrable. If $g : \Omega \rightarrow \overline{\mathbb{R}}$ is $\mathcal{G}$ -measurable and satisfies $\mathbb{E}[\mathbf{1}_C f] = \mathbb{E}[\mathbf{1}_C g]$ for all $C \in \mathcal{C}$, then g is a version of $\mathbb{E}_{\mathcal{G}}[\cdot, f]$.*

Proof. (i) We establish the forward direction by approximating h with simple functions, and deduce the converse from the uniqueness statement of Theorem 2.7.1.

Forward direction (non-negative f). Suppose first that $f : \Omega \rightarrow \overline{\mathbb{R}}_+$. When $h = \mathbf{1}_G$ for some $G \in \mathcal{G}$, the desired identity is exactly the defining property of $\mathbb{E}_{\mathcal{G}}[\cdot, f]$ (Eq. (2.10)). For a non-negative $\mathcal{G}$ -measurable simple function $h = \sum_{i=1}^n a_i \mathbf{1}_{G_i}$ with $a_i \in \mathbb{R}_+$ and $G_i \in \mathcal{G}$, the linearity of expectations (Lemma 2.5.2) extends this:

$$\mathbb{E}[hf] = \sum_{i=1}^n a_i \mathbb{E}[\mathbf{1}_{G_i} f] = \sum_{i=1}^n a_i \mathbb{E}[\mathbf{1}_{G_i} \mathbb{E}_{\mathcal{G}}[\cdot, f]] = \mathbb{E}[h\mathbb{E}_{\mathcal{G}}[\cdot, f]].$$

For an arbitrary $\mathcal{G}$ -measurable $h : \Omega \rightarrow \overline{\mathbb{R}}_+$, Theorem 2.2.6(ii) provides an increasing sequence $(h_n)_{n \in \mathbb{N}}$ of non-negative $\mathcal{G}$ -measurable simple functions with $h_n \uparrow h$ pointwise. Then $h_n f \uparrow hf$ and $h_n \mathbb{E}_{\mathcal{G}}[\cdot, f] \uparrow h \mathbb{E}_{\mathcal{G}}[\cdot, f]$ pointwise, so the monotone convergence theorem (Theorem 2.5.1), combined with the simple-function case applied to each h_n , yields

$$\mathbb{E}[hf] = \lim_{n \rightarrow \infty} \mathbb{E}[h_n f] = \lim_{n \rightarrow \infty} \mathbb{E}[h_n \mathbb{E}_{\mathcal{G}}[\cdot, f]] = \mathbb{E}[h \mathbb{E}_{\mathcal{G}}[\cdot, f]].$$

Forward direction (integrable f). Let $f : \Omega \rightarrow \overline{\mathbb{R}}$ be integrable, and decompose $f = f_+ - f_-$ into its non-negative integrable parts. Applying the non-negative case to each of f_+ and f_- gives

$$\mathbb{E}[hf_{\pm}] = \mathbb{E}[h \mathbb{E}_{\mathcal{G}}[\cdot, f_{\pm}]]$$

for every $\mathcal{G}$ -measurable $h : \Omega \rightarrow \overline{\mathbb{R}}_+$. Whenever hf is integrable—equivalently, both $\mathbb{E}[hf_+]$ and $\mathbb{E}[hf_-]$ are finite—all four expectations above are finite, and the linearity of expectations (Lemma 2.5.2) together with $\mathbb{E}_{\mathcal{G}}[\cdot, f] = \mathbb{E}_{\mathcal{G}}[\cdot, f_+] - \mathbb{E}_{\mathcal{G}}[\cdot, f_-]$ almost surely yields

$$\mathbb{E}[hf] = \mathbb{E}[hf_+] - \mathbb{E}[hf_-] = \mathbb{E}[h\mathbb{E}_{\mathcal{G}}[\cdot, f_+]] - \mathbb{E}[h\mathbb{E}_{\mathcal{G}}[\cdot, f_-]] = \mathbb{E}[h\mathbb{E}_{\mathcal{G}}[\cdot, f]].$$

Converse. Letting $h = \mathbf{1}_G$ for all $G \in \mathcal{G}$, we recover the defining property of the conditional expectation (Eq. (2.10)).

- (ii) We first note that $\mathbf{1}_A f$ and $\mathbf{1}_A g$ are integrable for every $A \in \mathcal{H}$: indeed, $(\mathbf{1}_A f)_{\pm} = \mathbf{1}_A f_{\pm} \leq f_{\pm}$ gives $\mathbb{E}[(\mathbf{1}_A f)_{\pm}] \leq \mathbb{E}[f_{\pm}] < \infty$ by Eq. (2.8), and similarly for g . In particular, by letting $A = \Omega$, g itself is integrable. Now define

$$\mathcal{D} = \{G \in \mathcal{G} : \mathbb{E}[\mathbf{1}_G f] = \mathbb{E}[\mathbf{1}_G g]\},$$

and let us check that $\mathcal{D}$ is a λ -system.

- $\Omega \in \mathcal{D}$, since $\Omega \in \mathcal{C}$ by hypothesis.
- Let $D, E \in \mathcal{D}$ with $E \subseteq D$, so that $D \setminus E \in \mathcal{G}$. Then $\mathbf{1}_{D \setminus E} f = \mathbf{1}_D f - \mathbf{1}_E f$ pointwise on Ω , and all the random variables involved are integrable, so the linearity of expectations (Lemma 2.5.2(ii)) gives

$$\mathbb{E}[\mathbf{1}_{D \setminus E} f] = \mathbb{E}[\mathbf{1}_D f] - \mathbb{E}[\mathbf{1}_E f] = \mathbb{E}[\mathbf{1}_D g] - \mathbb{E}[\mathbf{1}_E g] = \mathbb{E}[\mathbf{1}_{D \setminus E} g],$$

so $D \setminus E \in \mathcal{D}$.

- Let $D_n \in \mathcal{D}$ for all $n \in \mathbb{N}$ with $D_1 \subseteq D_2 \subseteq \dots$, and write $D = \cup_{n \in \mathbb{N}} D_n \in \mathcal{G}$. Then $\mathbf{1}_{D_n} f_{\pm} \uparrow \mathbf{1}_D f_{\pm}$ pointwise, so the monotone convergence theorem (Theorem 2.5.1) gives $\mathbb{E}[\mathbf{1}_{D_n} f_{\pm}] \rightarrow \mathbb{E}[\mathbf{1}_D f_{\pm}]$; as all these quantities are finite, $\mathbb{E}[\mathbf{1}_{D_n} f] \rightarrow \mathbb{E}[\mathbf{1}_D f]$, and the same argument applies to g . Hence

$$\mathbb{E}[\mathbf{1}_D f] = \lim_{n \rightarrow \infty} \mathbb{E}[\mathbf{1}_{D_n} f] = \lim_{n \rightarrow \infty} \mathbb{E}[\mathbf{1}_{D_n} g] = \mathbb{E}[\mathbf{1}_D g],$$

so $D \in \mathcal{D}$.

By hypothesis $\mathcal{C} \subseteq \mathcal{D}$, so the monotone class theorem (Theorem 2.1.5) yields $\mathcal{G} = \sigma(\mathcal{C}) \subseteq \mathcal{D}$; since $\mathcal{D} \subseteq \mathcal{G}$ by construction, $\mathcal{D} = \mathcal{G}$. In other words, $\mathbb{E}[\mathbf{1}_G f] = \mathbb{E}[\mathbf{1}_G g]$ for every $G \in \mathcal{G}$. □

The following result presents the conditional determinism property of conditional probabilities, which is of particular relevance to the theory of causal spaces.

Proposition 2.7.4 (Conditional determinism). *Let $\mathcal{G}$ be a sub- σ -algebra of $\mathcal{H}$, and $A \in \mathcal{H}$, $B \in \mathcal{G}$ events. For $\mathbb{P}$ -almost all $\gamma \in \Omega$, we have*

$$\mathbb{P}_{\mathcal{G}}(\gamma, A \cap B) = \mathbf{1}_B(\gamma) \mathbb{P}_{\mathcal{G}}(\gamma, A).$$

Proof. See that, for any $G \in \mathcal{G}$, by the defining property of conditional probabilities,

$$\mathbb{P}(A \cap B \cap G) = \int \mathbb{P}(d\gamma) \mathbf{1}_G(\gamma) \mathbb{P}_{\mathcal{G}}(\gamma, A \cap B) = \int \mathbb{P}(d\gamma) \mathbf{1}_G(\gamma) \mathbf{1}_B(\gamma) \mathbb{P}_{\mathcal{G}}(\gamma, A),$$

where, for the third expression, we used the fact that $B \cap G \in \mathcal{G}$. Now we can apply the almost sure uniqueness of conditional expectations (Theorem 2.7.1) to conclude that $\mathbb{P}_{\mathcal{G}}(\gamma, A \cap B) = \mathbf{1}_B(\gamma) \mathbb{P}_{\mathcal{G}}(\gamma, A)$ for $\mathbb{P}$ -almost all $\gamma \in \Omega$. $\square$

Now, the above definition of the conditional probability is defined for individual events $A \in \mathcal{H}$. Choosing a version of $\mathbb{P}_{\mathcal{G}}(\cdot, A)$ for each $A \in \mathcal{H}$, we have a map $\mathbb{P}_{\mathcal{G}} : \Omega \times \mathcal{H} \rightarrow [0, 1]$ given by $(\gamma, A) \mapsto \mathbb{P}_{\mathcal{G}}(\gamma, A)$. We say that $\mathbb{P}_{\mathcal{G}}$ is a *regular version* of the conditional probability if it is a transition probability kernel from $(\Omega, \mathcal{G})$ into $(\Omega, \mathcal{H})$.

We say that a finite number of events $A_1, \dots, A_n \in \mathcal{H}$ are *conditionally independent* given $\mathcal{G}$ if, for every non-empty subset $I \subseteq \{1, \dots, n\}$,

$$\mathbb{P}_{\mathcal{G}}(\gamma, \cap_{i \in I} A_i) = \prod_{i \in I} \mathbb{P}_{\mathcal{G}}(\gamma, A_i)$$

for $\mathbb{P}$ -almost all $\gamma \in \Omega$. We say that a finite number of sub- σ -algebras $\mathcal{H}_1, \dots, \mathcal{H}_n \subseteq \mathcal{H}$ are conditionally independent given $\mathcal{G}$ if all tuples of events $A_1 \in \mathcal{H}_1, \dots, A_n \in \mathcal{H}_n$ are conditionally independent given $\mathcal{G}$. Finally, an arbitrary set $\{\mathcal{H}_t\}_{t \in \mathbb{T}}$ of sub- σ -algebras of $\mathcal{H}$, where $\mathbb{T}$ can be uncountably infinite, are said to be conditionally independent given $\mathcal{G}$ if every finite subset of them are conditionally independent given $\mathcal{G}$.

The following result characterises (conditional) independence using conditional probabilities.

Proposition 2.7.5. *Two sub- σ -algebras $\mathcal{F}_1$ and $\mathcal{F}_2$ of $\mathcal{H}$ are independent if and only if, for all $A \in \mathcal{F}_2$, we have $\mathbb{P}_{\mathcal{F}_1}(\gamma, A) = \mathbb{P}(A)$ for $\mathbb{P}$ -almost all $\gamma \in \Omega$.*

More generally, $\mathcal{F}_1$ and $\mathcal{F}_2$ are conditionally independent given $\mathcal{G}$ if and only if, for all $A \in \mathcal{F}_2$, we have $\mathbb{P}_{\mathcal{G}}(\gamma, A) = \mathbb{P}_{\mathcal{F}_1 \vee \mathcal{G}}(\gamma, A)$ for $\mathbb{P}$ -almost all $\gamma \in \Omega$.

Proof. We prove the second statement; the first follows by taking $\mathcal{G} = \{\emptyset, \Omega\}$, since then $\mathbb{P}_{\mathcal{G}}(\cdot, A) = \mathbb{P}(A)$ and $\mathcal{F}_1 \vee \mathcal{G} = \mathcal{F}_1$.

($\implies$) Take any $B \in \mathcal{F}_1$, $A \in \mathcal{F}_2$ and $G \in \mathcal{G}$. By the conditional independence of $\mathcal{F}_1$ and $\mathcal{F}_2$ given $\mathcal{G}$, and applying Theorem 2.5.3(i),

$$\begin{aligned} \mathbb{E}[\mathbf{1}_{B \cap G} \mathbf{1}_A] &= \mathbb{E}[\mathbf{1}_G(\cdot) \mathbb{P}_{\mathcal{G}}(\cdot, A \cap B)] \\ &= \mathbb{E}[\mathbf{1}_G(\cdot) \mathbb{P}_{\mathcal{G}}(\cdot, A) \mathbb{P}_{\mathcal{G}}(\cdot, B)] \\ &= \mathbb{E}[\mathbf{1}_G(\cdot) \mathbb{P}_{\mathcal{G}}(\cdot, A) \mathbf{1}_B(\cdot)] \\ &= \mathbb{E}[\mathbf{1}_{B \cap G}(\cdot) \mathbb{P}_{\mathcal{G}}(\cdot, A)], \end{aligned}$$

where, for the third equality, we used Proposition 2.7.3(i) using the fact that $\mathbf{1}_G(\cdot) \mathbb{P}_{\mathcal{G}}(\cdot, A)$ is $\mathcal{G}$ -measurable. But the collection of intersections $B \cap G$ with $B \in \mathcal{F}_1$ and $G \in \mathcal{G}$ is a π -system that generates $\mathcal{F}_1 \vee \mathcal{G}$ (Lemma 2.1.3), and clearly contains Ω , so Proposition 2.7.3(ii) tells us that $\mathbb{P}_{\mathcal{G}}(\cdot, A)$ is a version of $\mathbb{P}_{\mathcal{F}_1 \vee \mathcal{G}}(\cdot, A)$.

($\Leftarrow$) Take any $B \in \mathcal{F}_1$, $A \in \mathcal{F}_2$ and $G \in \mathcal{G}$. Then

$$\begin{aligned}\mathbb{E}[\mathbf{1}_G(\cdot)\mathbb{P}_{\mathcal{G}}(\cdot, A \cap B)] &= \mathbb{E}[\mathbf{1}_G(\cdot)\mathbf{1}_B(\cdot)\mathbf{1}_A(\cdot)] \\ &= \mathbb{E}[\mathbf{1}_G(\cdot)\mathbf{1}_B(\cdot)\mathbb{P}_{\mathcal{F}_1 \vee \mathcal{G}}(\cdot, A)] \\ &= \mathbb{E}[\mathbf{1}_G(\cdot)\mathbf{1}_B(\cdot)\mathbb{P}_{\mathcal{G}}(\cdot, A)] \\ &= \mathbb{E}[\mathbf{1}_G(\cdot)\mathbb{P}_{\mathcal{G}}(\cdot, B)\mathbb{P}_{\mathcal{G}}(\cdot, A)],\end{aligned}$$

where, for the last equality, we used Proposition 2.7.3(i), together with the fact that $\mathbf{1}_G(\cdot)\mathbb{P}_{\mathcal{G}}(\cdot, A)$ is $\mathcal{G}$ -measurable. Now apply Theorem 2.5.3(ii) to conclude.

□

3 Warm-up: Homogeneous causal spaces

Probability theory, as reviewed briskly in the previous chapter, does not capture any causal information—the theory is only designed to encode the random occurrence of events, each with its own probability. Building upon the foundational objects and notions introduced in Chapter 2, we now begin our theory of causality, by introducing our main framework, namely, *causal spaces*. Recall that interventions are central to the notion of causality. A causal space is a mathematical structure which encodes information about the probabilities after interventions, and in which one can perform interventions.

When an intervention is carried out, many different events may occur, each with its own probability. To represent this inherent stochasticity, causal spaces are built on the objects and axioms of probability theory, reviewed in Chapter 2. What distinguishes a causal space from an ordinary probability space is an additional piece of structure, called a *causal mechanism*, which specifies how the probability measure changes when we intervene. The causal mechanism comes with its own *axioms*. Following the discussion on mathematical axiomatisations in Chapter 1, the axioms of the causal mechanism should capture, in a minimal way, what we believe to be an intervention. Informally, the axioms require that (i) doing nothing does not change anything, and (ii) when an intervention fixes a set of coordinates to a value, they indeed have this value after the intervention. This will be made precise in Section 3.1.

In general, when an intervention takes place, the effect of the intervention may be *heterogeneous* across subpopulations, depending on the pre-intervention information. For instance, the same medicine may have different effects depending on the patient’s age. In this chapter, we start by introducing *homogeneous causal spaces*, which only encode the causal effect after averaging out across the pre-intervention distribution across the entire population, and interventions are performed without regards to the pre-intervention values.

The theory of homogeneous causal spaces does not require meticulous care in measurability technicalities, yet gives a good intuition of what causal spaces are, and hence serves as a good warm-up before we introduce causal spaces in full generality, that take heterogeneity into account, in Chapter 4. Many remarks and discussions that apply to both cases will be given in this chapter, and the subsequent chapters will be reserved for material that is specific to general, heterogeneous causal spaces.

3.1 Homogeneous causal spaces

In order to define causal spaces, we require that the measurable space be in a product form (Section 2.1). The components in the product structure will play the role of *coordinates* or *variables* of the system. Let $\mathbb{T}$ be the index set of the product. Then taking,

for each $t \in \mathbb{T}$, a measurable space $(\Omega_t, \mathcal{E}_t)$, we have the product measurable space

$$(\Omega, \mathcal{H}) = \otimes_{t \in \mathbb{T}} (\Omega_t, \mathcal{E}_t) = (\times_{t \in \mathbb{T}} \Omega_t, \otimes_{t \in \mathbb{T}} \mathcal{E}_t),$$

as in Section 2.1. The semantic meaning of the coordinates and of the individual outcomes depends on the setting under consideration. For a study involving people, each outcome coordinate $\omega_t \in \Omega_t$ can be viewed as an attribute of a randomly drawn member of the population, for example, their height, weight, age, whether they take a medicine or blood pressure. In the study of stochastic processes, each outcome $\omega \in \Omega$ is a sample path of the process, and each $t \in \mathbb{T}$ a specific time point.

For studies involving people, we do not view Ω as literally the (finite) population of the world, nor each $\omega \in \Omega$ as an actual person from the population. Rather, each ω_t is interpreted as a possible value of that particular attribute of a single person, of which we only have access to a probability distribution. This is different to the approach typically taken in the potential outcomes framework (Section 1.4.1), in which each outcome is viewed as a “unit” of the experiment and the potential outcomes are *fixed* attributes of those units. If we want to model multiple people explicitly, then we copy the entire measurable space the corresponding number of times, and model the coordinate values of each person as being random.

Recall also from Section 2.1 that, for $\mathbb{S} \in \mathcal{P}(\mathbb{T})$, we write $\mathcal{H}_{\mathbb{S}}$ for the sub- σ -algebra of $\mathcal{H}$ corresponding to the index set $\mathbb{S}$. Further, each outcome $\omega \in \Omega$ decomposes as $\omega = (\omega_t)_{t \in \mathbb{T}}$, and, for each $\mathbb{S} \in \mathcal{P}(\mathbb{T})$, we write $\omega_{\mathbb{S}} = (\omega_s)_{s \in \mathbb{S}} \in \Omega_{\mathbb{S}} = \times_{s \in \mathbb{S}} \Omega_s$.

We are ready to introduce the definition of homogeneous causal spaces. The concepts of product measurable spaces and transition probability kernels introduced in Chapter 2 are crucial.

Definition 3.1.1 (Homogeneous causal spaces). A *homogeneous causal space* $(\Omega, \mathcal{H}, \mathbb{P}, \mathbb{K})$ consists of

- a product measurable space $(\Omega, \mathcal{H}) = \otimes_{t \in \mathbb{T}} (\Omega_t, \mathcal{E}_t)$;
- a probability measure $\mathbb{P}$ on $(\Omega, \mathcal{H})$; and
- a *causal mechanism* $\mathbb{K} = \{K_{\mathbb{S}} : \mathbb{S} \in \mathcal{P}(\mathbb{T})\}$, which is a collection of transition probability kernels $K_{\mathbb{S}}$, called *causal kernels*, from $(\Omega, \mathcal{H}_{\mathbb{S}})$ into $(\Omega, \mathcal{H})$.

The causal kernels $K_{\mathbb{S}}$ are required to satisfy the following axioms:

- (i) *trivial intervention*: for all $A \in \mathcal{H}$ and $\omega \in \Omega$,

$$K_{\emptyset}(\omega; A) = \mathbb{P}(A);$$

- (ii) *interventional determinism*: for all $\mathbb{S} \in \mathcal{P}(\mathbb{T})$, $B \in \mathcal{H}_{\mathbb{S}}$ and $\omega \in \Omega$,

$$K_{\mathbb{S}}(\omega; B) = \mathbf{1}_B(\omega) = \delta_{\omega}(B).$$




Figure 3.1: The causal kernel $K_{\mathbb{U}}(\omega; \cdot)$ (Definition 3.1.1) is the measure on $\mathcal{H}$ when fixing the coordinates $\mathbb{U}$ to the value $\omega_{\mathbb{U}}$. The causal kernel satisfies the interventional determinism axiom (Axiom (ii)), so the measure on $\mathcal{H}_{\mathbb{U}}$ is the delta measure δ_{ω} .

In words, we have a quadruple $(\Omega, \mathcal{H}, \mathbb{P}, \mathbb{K})$, where the first triple $(\Omega, \mathcal{H}, \mathbb{P})$ forms a probability space as introduced in Chapter 2, with the measurable space additionally required to be in the form of a product. The causal mechanism $\mathbb{K}$ is the additional object not present in probability spaces.

The high level idea is as follows. The measure $\mathbb{P}$ should be thought of as the *observational* measure, i.e. the probability measure that the space is endowed with before any intervention takes place. Thus, $\mathbb{P}$ contains no *causal* information whatsoever (at least without additional assumptions—see Section 3.4). The causal mechanism $\mathbb{K}$ is where the causal, or *interventional*, information is encoded. Specifically, for a set $\mathbb{U} \in \mathcal{P}(\mathbb{T})$, the corresponding causal kernel $K_{\mathbb{U}}$ tells us what the measure on the space will be after fixing the $\mathbb{U}$ -coordinates. See Fig. 3.1 for an illustration. The precise definition of an intervention in homogeneous causal spaces will be given in Section 3.2.

The causal kernels come with two axioms, namely, the *trivial intervention* axiom (Axiom (i)) and the *interventional determinism* axiom (Axiom (ii)). They capture, in a minimal way, our intuition of an intervention:

- Axiom (i) says that intervening on the empty set $\emptyset$ (i.e. doing nothing) keeps the observational measure $\mathbb{P}$ intact; in other words, the measure on the space $(\Omega, \mathcal{H})$ does not change.
- Axiom (ii) says that, after fixing the $\mathbb{U}$ -coordinates to $\omega_{\mathbb{U}}$, the $\mathbb{U}$ -coordinates indeed have those values $\omega_{\mathbb{U}}$. The causal kernel assigns zero measure to events that do not allow the $\mathbb{U}$ -coordinates to be $\omega_{\mathbb{U}}$.

These are the non-negotiable properties that any causal kernel must satisfy, and if these are violated, we no longer believe that this mathematical object captures what we mean

by an intervention. On the other hand, no other restrictions or modelling assumptions are placed on the causal kernels—the axioms capture what we mean by an intervention in a minimal and general way.

Let us investigate the interventional determinism axiom in more depth. The term *determinism* in Axiom (ii) refers to the coordinates $\mathbb{U}$ that are intervened on; the remaining coordinates may still be random, with the measure captured by the corresponding causal kernel $K_{\mathbb{U}}(\omega; \cdot)$. We have the following result that extends Axiom (ii) to all events in $\mathcal{H}$, beyond $\mathcal{H}_{\mathbb{U}}$, as well as measurable functions.

Theorem 3.1.2 (Interventional determinism). *Take any $\mathbb{U} \in \mathcal{P}(\mathbb{T})$. The interventional determinism axiom (Axiom (ii)) is equivalent to the following three statements.*

(a) For all $A \in \mathcal{H}$, all $B \in \mathcal{H}_{\mathbb{U}}$ and all $\omega \in \Omega$,

$$K_{\mathbb{U}}(\omega; A \cap B) = \mathbf{1}_B(\omega) K_{\mathbb{U}}(\omega; A).$$

(b) If $f : \Omega \rightarrow \overline{\mathbb{R}}_+$ is $\mathcal{H}$ -measurable, or if $f : \Omega \rightarrow \overline{\mathbb{R}}$ is $\mathcal{H}$ -measurable and integrable, then for all $\omega \in \Omega$,

$$\int K_{\mathbb{U}}(\omega; d\omega') f(\omega') = \int K_{\mathbb{U}}(\omega; d\omega') f((\omega_{\mathbb{U}}, \omega'_{\mathbb{T} \setminus \mathbb{U}})).$$

(c) Suppose that $g, h : \Omega \rightarrow \overline{\mathbb{R}}$ are $\mathcal{H}_{\mathbb{U}}$ - and $\mathcal{H}$ -measurable respectively. Suppose further that the following two conditions are satisfied.

- Either $gh : \Omega \rightarrow \overline{\mathbb{R}}_+$ or $gh : \Omega \rightarrow \overline{\mathbb{R}}$ is integrable.
- Either $h : \Omega \rightarrow \overline{\mathbb{R}}_+$ or $h : \Omega \rightarrow \overline{\mathbb{R}}$ is integrable.

Then for all $\omega \in \Omega$,

$$\int K_{\mathbb{U}}(\omega; d\omega') g(\omega') h(\omega') = g(\omega) \int K_{\mathbb{U}}(\omega; d\omega') h(\omega').$$

Proof. **Axiom (ii) $\implies$ (a)** We look at the following two cases separately.

- $\omega \notin B$: We have $\mathbf{1}_B(\omega) = 0$; see that, by Axiom (ii),

$$K_{\mathbb{U}}(\omega; A \cap B) \leq K_{\mathbb{U}}(\omega; B) = \mathbf{1}_B(\omega) = 0,$$

giving us $K_{\mathbb{U}}(\omega; A \cap B) = 0 = \mathbf{1}_B(\omega) K_{\mathbb{U}}(\omega; A)$.

- $\omega \in B$: We have $K_{\mathbb{U}}(\omega; B) = \mathbf{1}_B(\omega) = 1$ by Axiom (ii), i.e. B has probability 1 under the measure $K_{\mathbb{U}}(\omega; \cdot)$. Hence, for all events $A \in \mathcal{H}$,

$$K_{\mathbb{U}}(\omega; A \cap B) = K_{\mathbb{U}}(\omega; A) = \mathbf{1}_B(\omega) K_{\mathbb{U}}(\omega; A),$$

as required.

(a) $\implies$ (b) Recall the canonical projection and insertion maps from Section 2.2. For each $\omega \in \Omega$, the map $\iota_{\omega_{\mathbb{U}}} \circ \pi_{\mathbb{T} \setminus \mathbb{U}} : \Omega \rightarrow \Omega$ overwrites the $\mathbb{U}$ -components with $\omega_{\mathbb{U}}$, i.e. by $\iota_{\omega_{\mathbb{U}}} \circ \pi_{\mathbb{T} \setminus \mathbb{U}}(\omega') = (\omega_{\mathbb{U}}, \omega'_{\mathbb{T} \setminus \mathbb{U}})$. It is measurable, since it is a composition of measurable functions. Now, for any measurable cylinders $B \cap C$ with $B \in \mathcal{H}_{\mathbb{U}}$ and $C \in \mathcal{H}_{\mathbb{T} \setminus \mathbb{U}}$, part (a) gives

$$K_{\mathbb{U}}(\omega; B \cap C) = \mathbf{1}_B(\omega) K_{\mathbb{U}}(\omega; C).$$

On the other hand, see that

$$(\iota_{\omega_{\mathbb{U}}} \circ \pi_{\mathbb{T} \setminus \mathbb{U}})^{-1}(B \cap C) = \{\omega' \in \Omega : (\omega_{\mathbb{U}}, \omega'_{\mathbb{T} \setminus \mathbb{U}}) \in B \cap C\} = \begin{cases} C & \text{if } \omega \in B \\ \emptyset & \text{otherwise,} \end{cases}$$

hence we also have

$$K_{\mathbb{U}}(\omega; (\iota_{\omega_{\mathbb{U}}} \circ \pi_{\mathbb{T} \setminus \mathbb{U}})^{-1}(B \cap C)) = \begin{cases} K_{\mathbb{U}}(\omega; C) & \text{if } \omega \in B \\ 0 & \text{otherwise} \end{cases} = \mathbf{1}_B(\omega) K_{\mathbb{U}}(\omega; C),$$

which is the same as $K_{\mathbb{U}}(\omega; B \cap C)$ by (a). So the two measures $K_{\mathbb{U}}(\omega; \cdot)$ and $K_{\mathbb{U}}(\omega; (\iota_{\omega_{\mathbb{U}}} \circ \pi_{\mathbb{T} \setminus \mathbb{U}})^{-1}(\cdot))$ coincide on the semi-ring of measurable cylinders, and by the unique extension to all of $\mathcal{H}$ given by the Carathéodory extension theorem (Theorem 2.4.1), these two measures coincide on all of $\mathcal{H}$.

Now, take any $\mathcal{H}$ -measurable function $f : \Omega \rightarrow \overline{\mathbb{R}}$. Then by the above argument and by applying the formula for integration with respect to image measures (Lemma 2.5.7),

$$\begin{aligned} \int K_{\mathbb{U}}(\omega; d\omega') f(\omega') &= \int K_{\mathbb{U}}(\omega; (\iota_{\omega_{\mathbb{U}}} \circ \pi_{\mathbb{T} \setminus \mathbb{U}})^{-1}(d\omega')) f(\omega') \\ &= \int K_{\mathbb{U}}(\omega; d\omega') f(\iota_{\omega_{\mathbb{U}}} \circ \pi_{\mathbb{T} \setminus \mathbb{U}}(\omega')) \\ &= \int K_{\mathbb{U}}(\omega; d\omega') f((\omega_{\mathbb{U}}, \omega'_{\mathbb{T} \setminus \mathbb{U}})), \end{aligned}$$

which is precisely (b).

(b) $\implies$ (c) In (b), let $f = gh$. Then

$$\begin{aligned} \int K_{\mathbb{U}}(\omega; d\omega') g(\omega') h(\omega') &= \int K_{\mathbb{U}}(\omega; d\omega') g((\omega_{\mathbb{U}}, \omega'_{\mathbb{T} \setminus \mathbb{U}})) h((\omega_{\mathbb{U}}, \omega'_{\mathbb{T} \setminus \mathbb{U}})) \\ &= g(\omega) \int K_{\mathbb{U}}(\omega; d\omega') h((\omega_{\mathbb{U}}, \omega'_{\mathbb{T} \setminus \mathbb{U}})) \quad \text{since } g \text{ is } \mathcal{H}_{\mathbb{U}}\text{-measurable} \\ &= g(\omega) \int K_{\mathbb{U}}(\omega; d\omega') h(\omega'), \end{aligned}$$

where we applied part (b) again on the last line with $f = h$. This is precisely (c).

(c) $\implies$ **Axiom (ii)** For any $B \in \mathcal{H}_{\mathbb{U}}$, let $g = \mathbf{1}_B$ and $h = 1$. Then (c) immediately tells us that

$$K_{\mathbb{U}}(\omega; B) = \int K_{\mathbb{U}}(\omega; d\omega') \mathbf{1}_B(\omega') = \mathbf{1}_B(\omega),$$

as required. $\square$

Note also that the equivalent form of the interventional determinism axiom given in part (a) is precisely analogous to the conditional determinism property of conditional probability measures as proved in Proposition 2.7.4, with the conditional probability replaced by a causal kernel. The difference lies in what the primitive object is, and whether determinism is a *property* or an *axiom*.

- In probability theory, the measure $\mathbb{P}$ is the primitive object, the conditional measure $\mathbb{P}_{\mathcal{G}}$ is derived from $\mathbb{P}$, and conditional determinism is a property of $\mathbb{P}_{\mathcal{G}}$ that can be proved using the definition of $\mathbb{P}_{\mathcal{G}}$.
- In causal spaces, the causal kernels $K_{\mathbb{S}}$ are themselves primitive objects, and interventional determinism is a defining axiom of the causal kernels. General transition probability kernels do not necessarily possess this property.

The fact that the causal kernels are the *primitive objects* of the space, in addition to the probability measure $\mathbb{P}$, is important. A priori, the causal mechanism $\mathbb{K}$ is not related to the measure $\mathbb{P}$ in any way, unless additional assumptions are imposed (see Section 3.4). This is what sets them apart from *conditional probabilities*, which are fully determined once $\mathbb{P}$ is specified. Indeed, the fact that intervening and conditioning are different is the main premise of the theory of causality.

The following remark will be very important throughout the monograph.

Remark 3.1.3 (“Almost all” vs. “all” outcomes). It is important to note that the causal kernels $K_{\mathbb{S}}(\omega; \cdot)$ in a causal space (Definition 3.1.1) are defined *for all outcomes* $\omega \in \Omega$. This differs markedly from probability theory, where one is rarely interested in events of zero measure, and notions such as conditional expectations, conditional probabilities and conditional independence (Section 2.7) are only defined up to $\mathbb{P}$ -null events. However, in causal spaces, it is necessary that causal kernels are defined *for all* $\omega \in \Omega$, since the $\mathbb{S}$ -coordinates may be fixed to any value $\omega_{\mathbb{S}}$ during an intervention, and we want the subsequent measure to be well-defined. Throughout the monograph, we will repeatedly discuss the interplay between probabilistic and causal notions, and it is crucial to keep track of what is defined for almost all outcomes (with respect to a particular measure), and what is defined for all outcomes.

We now present some consequences of the interventional determinism axiom when we are conditioning the causal kernel $K_{\mathbb{U}}(\omega; \cdot)$ on a sub- σ -algebra $\mathcal{G}$. Here, the care in dealing with almost all outcomes and all outcomes, as discussed in Remark 3.1.3, is particularly pertinent. Purely for notational convenience, we write $\mathbb{P}^{\text{do}(\mathbb{U}, \delta_{\omega})}(\cdot)$ for the measure $K_{\mathbb{U}}(\omega; \cdot)$ for each $\omega \in \Omega$, since conditioning on this measure will add subscripts. This notation is for *interventional probabilities*, that will be formally introduced in Section 3.2.

Lemma 3.1.4. Take some $\mathbb{U} \in \mathcal{P}(\mathbb{T})$, a sub- σ -algebra $\mathcal{G} \subseteq \mathcal{H}$ and an outcome $\omega \in \Omega$.

- (i) For any $A \in \mathcal{H}$, we have $\mathbb{P}_{\mathcal{H}_{\mathbb{U}} \vee \mathcal{G}}^{\text{do}(\mathbb{U}, \delta_{\omega})}(\gamma, A) = \mathbb{P}_{\mathcal{G}}^{\text{do}(\mathbb{U}, \delta_{\omega})}(\gamma, A)$ for $\mathbb{P}^{\text{do}(\mathbb{U}, \delta_{\omega})}$ -almost all $\gamma \in \Omega$.
- (ii) For any $B \in \mathcal{H}_{\mathbb{U}}$, we have $\mathbb{P}_{\mathcal{G}}^{\text{do}(\mathbb{U}, \delta_{\omega})}(\gamma, B) = \mathbf{1}_B(\omega)$ for $\mathbb{P}^{\text{do}(\mathbb{U}, \delta_{\omega})}$ -almost all $\gamma \in \Omega$.
- (iii) For any $A \in \mathcal{H}$ and $B \in \mathcal{H}_{\mathbb{U}}$, we have $\mathbb{P}_{\mathcal{G}}^{\text{do}(\mathbb{U}, \delta_{\omega})}(\gamma, A \cap B) = \mathbf{1}_B(\omega) \mathbb{P}_{\mathcal{G}}^{\text{do}(\mathbb{U}, \delta_{\omega})}(\gamma, A)$ for $\mathbb{P}^{\text{do}(\mathbb{U}, \delta_{\omega})}$ -almost all $\gamma \in \Omega$.

Proof. (i) For any $B \in \mathcal{H}_{\mathbb{U}}$ and $G \in \mathcal{G}$, we can apply Theorem 3.1.2(c) to move the $\mathbf{1}_B$ in and out of the integral as follows:

$$\begin{aligned}
 \int K_{\mathbb{U}}(\omega; d\gamma) \mathbf{1}_B(\gamma) \mathbf{1}_G(\gamma) \mathbb{P}_{\mathcal{H}_{\mathbb{U}} \vee \mathcal{G}}^{\text{do}(\mathbb{U}, \delta_{\omega})}(\gamma, A) &= \mathbb{P}^{\text{do}(\mathbb{U}, \delta_{\omega})}(A \cap B \cap G) \\
 &= \mathbf{1}_B(\omega) \int K_{\mathbb{U}}(\omega; d\gamma) \mathbf{1}_G(\gamma) \mathbb{P}_{\mathcal{H}_{\mathbb{U}} \vee \mathcal{G}}^{\text{do}(\mathbb{U}, \delta_{\omega})}(\gamma, A) \\
 &= \mathbf{1}_B(\omega) \mathbb{P}^{\text{do}(\mathbb{U}, \delta_{\omega})}(A \cap G) \\
 &= \mathbf{1}_B(\omega) \int K_{\mathbb{U}}(\omega; d\gamma) \mathbf{1}_G(\gamma) \mathbb{P}_{\mathcal{G}}^{\text{do}(\mathbb{U}, \delta_{\omega})}(\gamma, A) \\
 &= \int K_{\mathbb{U}}(\omega; d\gamma) \mathbf{1}_B(\gamma) \mathbf{1}_G(\gamma) \mathbb{P}_{\mathcal{G}}^{\text{do}(\mathbb{U}, \delta_{\omega})}(\gamma, A).
 \end{aligned}$$

Now apply Proposition 2.7.3(ii) to conclude.

- (ii) The constant function $\gamma \mapsto \mathbf{1}_B(\omega)$ is trivially $\mathcal{G}$ -measurable, and Theorem 3.1.2(a) gives, for every $G \in \mathcal{G}$,

$$\mathbb{P}^{\text{do}(\mathbb{U}, \delta_{\omega})}(B \cap G) = K_{\mathbb{U}}(\omega; B \cap G) = \mathbf{1}_B(\omega) K_{\mathbb{U}}(\omega; G) = \mathbb{E}^{\text{do}(\mathbb{U}, \delta_{\omega})}[\mathbf{1}_G(\cdot) \mathbf{1}_B(\omega)],$$

so that the constant $\mathbf{1}_B(\omega)$ is a version of the conditional probability of B given $\mathcal{G}$ under $\mathbb{P}^{\text{do}(\mathbb{U}, \delta_{\omega})}$.

- (iii) See that, for any $G \in \mathcal{G}$,

$$\begin{aligned}
 \mathbb{E}^{\text{do}(\mathbb{U}, \delta_{\omega})}[\mathbf{1}_G(\cdot) \mathbb{P}_{\mathcal{G}}^{\text{do}(\mathbb{U}, \delta_{\omega})}(\cdot, A \cap B)] &= \mathbb{P}^{\text{do}(\mathbb{U}, \delta_{\omega})}(A \cap B \cap G) \\
 &= K_{\mathbb{U}}(\omega; A \cap B \cap G) \\
 &= \mathbf{1}_B(\omega) K_{\mathbb{U}}(\omega; A \cap G) \\
 &= \mathbb{E}^{\text{do}(\mathbb{U}, \delta_{\omega})}[\mathbf{1}_G(\cdot) \mathbf{1}_B(\omega) \mathbb{P}_{\mathcal{G}}^{\text{do}(\mathbb{U}, \delta_{\omega})}(\cdot, A)],
 \end{aligned}$$

hence, by Theorem 2.5.3(ii), we have $\mathbb{P}_{\mathcal{G}}^{\text{do}(\mathbb{U}, \delta_{\omega})}(\gamma, A \cap B) = \mathbf{1}_B(\omega) \mathbb{P}_{\mathcal{G}}^{\text{do}(\mathbb{U}, \delta_{\omega})}(\gamma, A)$ for $\mathbb{P}^{\text{do}(\mathbb{U}, \delta_{\omega})}$ -almost all $\gamma \in \Omega$. □

Let us give some concrete instantiations of (homogeneous) causal spaces.

Outcome	$\omega = (\omega_1, \omega_2)$			
	(0, 0)	(0, 1)	(1, 0)	(1, 1)
$\mathbb{P}$	0.4	0.1	0.1	0.4
$K_1(0; \cdot)$	0.5	0.5	0	0
$K_1(1; \cdot)$	0	0	0.5	0.5
$K_2(0; \cdot)$	0.5	0	0.5	0
$K_2(1; \cdot)$	0	0.5	0	0.5

Table 3.1: In Example 3.1.5, $\mathcal{H}_1$ and $\mathcal{H}_2$ are dependent, but there is no causation between them.

Outcome	$\omega = (\omega_{\text{Med}}, \omega_{\text{Blood}})$			
	(0, 0)	(0, 1)	(1, 0)	(1, 1)
$\mathbb{P}$	0.9	0.06	0.02	0.02
$K_{\text{Med}}(0; \cdot)$	0.93	0.07	0	0
$K_{\text{Med}}(1; \cdot)$	0	0	0.97	0.03

Table 3.2: In Example 3.1.6, making the patient (not) take the medicine changes the measure on their blood pressure.

Example 3.1.5 (Correlation does not imply causation). Let us take $\Omega_1 = \Omega_2 = \{0, 1\}$, and $\Omega = \Omega_1 \times \Omega_2$, with the respective power set σ -algebras. The observational probability measure and the causal kernels are specified in Table 3.1 ($K_{\{1,2\}}$ is determined by interventional determinism). We can see that the marginal observational measure is

$$\mathbb{P}(\omega_1 = 0) = \mathbb{P}(\omega_1 = 1) = \frac{1}{2} = \mathbb{P}(\omega_2 = 0) = \mathbb{P}(\omega_2 = 1),$$

but the conditionals are $\mathbb{P}_{\omega_2=0}(\omega_1 = 0) = 0.8$ and $\mathbb{P}_{\omega_2=1}(\omega_1 = 0) = 0.2$, so under the observational measure $\mathbb{P}$, the two coordinate σ -algebras $\mathcal{H}_1$ and $\mathcal{H}_2$ are clearly dependent. However, the causal kernels tell us that all the marginals remain at $\frac{1}{2}$, meaning that there is no causation between the two coordinates. This represents the classic mantra that “correlation does not imply causation”, and we stress that this cannot be captured by probability theory alone, where the observational measure $\mathbb{P}$ is all that is imposed on the space. $\square$

In the next example, there is causation between the two variables, in addition to dependence.

Example 3.1.6. Let us model patients taking a medicine and their blood pressure, with the outcome sets $\Omega_{\text{Med}} = \Omega_{\text{Blood}} = \{0, 1\}$. Semantically, $\omega_{\text{Med}} = 0$ means the patient does not take the medicine, $\omega_{\text{Med}} = 1$ means the patient takes the medicine, $\omega_{\text{Blood}} = 0$ means the blood pressure is normal and $\omega_{\text{Blood}} = 1$ means the blood pressure is not normal. Again, we take the power set σ -algebras, and the probability measure and the causal kernels are specified in Table 3.2.

As the $\mathbb{P}$ row tells us, only 4% of the population takes the medicine of their own accord, but among those who do, it is equally likely for the blood pressure to be in the normal range, as to be outside. In the whole population, the marginal probability of the

blood pressure being normal is 0.92, obtained by adding the first and third columns of the first row.

If the patient is told not to take the medicine (second row, $K_{\text{Med}}(0; \cdot)$), then the blood pressure of the patient will be in the normal range with 93% probability. On the other hand, if the patient is given the medicine, the probability that their blood pressure will be normal is 97%. $\square$

Recall that Definition 3.1.1 requires a product measurable space structure as the underlying measurable space in a causal space. Let us end this section by explaining the importance of this product structure in the following remark.

Remark 3.1.7 (Product structure, I). The required product measurable space structure $(\Omega, \mathcal{H}) = \otimes_{t \in \mathbb{T}} (\Omega_t, \mathcal{E}_t)$ with index set $\mathbb{T}$ is an essential part of the causal space specification, and the causal mechanism $\mathbb{K}$ explicitly depends on this product structure. This stands in contrast to probability theory, where the product structure of the measurable space, if it is present, is merely a convenient (albeit ubiquitous) way of constructing a measurable space.

More concretely, let us consider sets $\Omega_1 = \Omega_2 = \{1, 2\}$ and $\Omega_3 = \{1, 2, 3, 4\}$, all equipped with the power-set σ -algebras $\mathcal{E}_1, \mathcal{E}_2, \mathcal{E}_3$ respectively. Then clearly, as probability spaces, $(\Omega_1 \times \Omega_2, \mathcal{E}_1 \otimes \mathcal{E}_2, \mathbb{P})$ and $(\Omega_3, \mathcal{E}_3, \mathbb{P})$ are isomorphic, for any measure $\mathbb{P}$ that assigns the same measure to the corresponding elements. However, if we are to turn these into causal spaces, the product structure explicitly enters into the construction. For the former, the causal mechanism $\mathbb{K}$ will consist of the causal kernels $K_\emptyset, K_{\{1\}}, K_{\{2\}}, K_{\{1,2\}}$, respectively corresponding to interventions on nothing, on each component of the product, and on both components. On the other hand, ignoring the product structure as in $(\Omega_3, \mathcal{E}_3, \mathbb{P})$, the causal mechanism $\mathbb{K}'$ will only consist of $K_\emptyset$ and K_3 . Then the resulting causal spaces $(\Omega_1 \times \Omega_2, \mathcal{E}_1 \otimes \mathcal{E}_2, \mathbb{P}, \mathbb{K})$ and $(\Omega_3, \mathcal{E}_3, \mathbb{P}, \mathbb{K}')$ will be structurally distinct.

3.2 Interventions on homogeneous causal spaces

Now that we have defined homogeneous causal spaces, we move onto the central operation of interventions. Recall that each causal kernel $K_{\mathbb{S}}(\omega; \cdot)$ is the measure on the whole space $(\Omega, \mathcal{H})$ when fixing the $\mathbb{S}$ -coordinate to $\omega_{\mathbb{S}}$. In general, we will place a *measure* on the σ -algebra $\mathcal{H}_{\mathbb{S}}$ rather than fixing it to a single value.

In full generality, an intervention can also be heterogeneous, where different measures can be imposed depending on the pre-intervention values. An intervention can even be *targeted*, whereby the coordinates that are intervened on themselves depend on the pre-intervention values. These general forms of interventions will be introduced in full in Chapter 5, but in this warm-up chapter, we only introduce homogeneous interventions, which are the only type of interventions that homogeneous causal spaces can carry. This greatly simplifies the measure-theoretic technicalities.

Definition 3.2.1 (Intervention). Let $\mathbb{Q}$ be a probability measure on $(\Omega, \mathcal{H}_{\mathbb{U}})$, called the *assignment measure*. An *intervention* on $\mathcal{H}_{\mathbb{U}}$ via $\mathbb{Q}$ yields a new, *interventional*




Figure 3.2: The interventional probability measure $\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}$ after an intervention on the coordinates $\mathcal{H}_U$ with an assignment measure $\mathbb{Q}$ (Definition 3.2.1 and Eq. (3.1)). The measure $\mathbb{Q}$ on $\mathcal{H}_U$ is propagated to all of $\mathcal{H}$ via K_U to yield $\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})} = \int \mathbb{Q}(d\omega) K_U(\omega; \cdot)$.

causal space $(\Omega, \mathcal{H}, \mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}, \mathbb{K}^{\text{do}(\mathbb{U}, \mathbb{Q})})$, where

- the *interventional probability measure* $\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}$ on $(\Omega, \mathcal{H})$ is defined, for $A \in \mathcal{H}$, by

$$\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}(A) = \int \mathbb{Q}(d\omega) K_U(\omega; A); \quad (3.1)$$

- $\mathbb{K}^{\text{do}(\mathbb{U}, \mathbb{Q})} = \{K_S^{\text{do}(\mathbb{U}, \mathbb{Q})} : S \in \mathcal{P}(\mathbb{T})\}$ is the *interventional causal mechanism*, where the *interventional causal kernels* $K_S^{\text{do}(\mathbb{U}, \mathbb{Q})}$ are transition probability kernels from $(\Omega, \mathcal{H}_S)$ into $(\Omega, \mathcal{H})$, given, for $A \in \mathcal{H}$, by

$$K_S^{\text{do}(\mathbb{U}, \mathbb{Q})}(\omega; A) = \int \mathbb{Q}(d\omega'_{U \setminus S}) K_{S \cup U}((\omega_S, \omega'_{U \setminus S}); A). \quad (3.2)$$

Let us go through the intuition behind this definition. First, when the probability on $\mathcal{H}_U$ is changed from the marginal observational measure $\mathbb{P}$ to an arbitrary assignment measure $\mathbb{Q}$, the corresponding causal kernel K_U propagates $\mathbb{Q}$ to the rest of the measurable space $(\Omega, \mathcal{H})$, to yield the interventional probability measure $\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}$ in Eq. (3.1). See Fig. 3.2 for an illustration.

Let us investigate what the two axioms of (homogeneous) causal spaces (Definition 3.1.1) signify in light of the definition of interventions.

- When we intervene on the trivial σ -algebra $(\Omega, \mathcal{H}_\emptyset)$, the only possible assignment measure $\mathbb{Q}$ is such that $\mathbb{Q}(\emptyset) = 0$ and $\mathbb{Q}(\Omega) = 1$. Then, according to Eq. (3.1), the

interventional probability measure is calculated as

$$\mathbb{P}^{\text{do}(\emptyset, \mathbb{Q})}(A) = \int \mathbb{Q}(d\omega) K_{\emptyset}(\omega; A) = \int \mathbb{Q}(d\omega) \mathbb{P}(A) = \mathbb{P}(A)$$

for any $A \in \mathcal{H}$, where we applied Axiom (i). Hence, Axiom (i) ensures that the interventional measure $\mathbb{P}^{\text{do}(\emptyset, \mathbb{Q})}$ remains the same as the observational measure $\mathbb{P}$ for all events $A \in \mathcal{H}$.

- (ii) Let us take any $B \in \mathcal{H}_{\mathbb{U}}$. Then, according to Eq. (3.1), the interventional probability measure is calculated as

$$\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}(B) = \int \mathbb{Q}(d\omega) K_{\mathbb{U}}(\omega; B) = \int \mathbb{Q}(d\omega) \mathbf{1}_B(\omega) = \mathbb{Q}(B),$$

where we applied Axiom (ii). Hence, Axiom (ii) ensures that the interventional measure $\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}$ coincides with the assignment measure $\mathbb{Q}$ on $\mathcal{H}_{\mathbb{U}}$ where the intervention took place.

More generally, the following lemma gives an integral form of the above consequence of the determinism axiom.

Lemma 3.2.2. *For an $\mathcal{H}_{\mathbb{U}}$ -measurable function $f : \Omega \rightarrow \overline{\mathbb{R}}_+$, or an $\mathcal{H}_{\mathbb{U}}$ -measurable and integrable function $f : \Omega \rightarrow \overline{\mathbb{R}}$, we have*

$$\int \mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}(d\omega) f(\omega) = \int \mathbb{Q}(d\omega) f(\omega).$$

Proof. See that, by applying Theorem 3.1.2(c), we immediately have

$$\int \mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}(d\omega) f(\omega) = \int \int \mathbb{Q}(d\omega') K_{\mathbb{U}}(\omega'; d\omega) f(\omega) = \int \mathbb{Q}(d\omega') f(\omega').$$

□

Let us now turn our attention to the interventional causal kernels (Eq. (3.2)). Each $K_{\mathbb{S}}^{\text{do}(\mathbb{U}, \mathbb{Q})}(\omega; \cdot)$ is the measure that will be obtained on the whole space $(\Omega, \mathcal{H})$ when fixing the $\mathbb{S}$ -coordinates to ω in the interventional causal space $(\Omega, \mathcal{H}, \mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}, \mathbb{K}^{\text{do}(\mathbb{U}, \mathbb{Q})})$. In other words, the interventional causal kernels tell us about *sequential* interventions, first on $\mathcal{H}_{\mathbb{U}}$, then on $\mathcal{H}_{\mathbb{S}}$. The $\mathbb{S}$ -coordinates (where the second intervention takes place) will be fixed to $\omega_{\mathbb{S}}$, and the $(\mathbb{U} \setminus \mathbb{S})$ -coordinates that received the assignment measure $\mathbb{Q}$ in the first intervention but were not intervened on in the second intervention will retain the measure $\mathbb{Q}$. In particular, if $\mathbb{S} \cap \mathbb{U} = \emptyset$, then $\mathcal{H}_{\mathbb{U}}$ retains $\mathbb{Q}$ in its entirety, whereas if $\mathbb{S} \supset \mathbb{U}$, then the first assignment $\mathbb{Q}$ is overwritten completely. Finally, the product measure $\delta_{\omega|_{\mathbb{S}}} \otimes \mathbb{Q}|_{\mathbb{U} \setminus \mathbb{S}}$ on $\mathcal{H}_{\mathbb{S} \cup \mathbb{U}}$ is propagated to the rest of the space via $K_{\mathbb{S} \cup \mathbb{U}}$.

It is necessary to show that the interventional causal space $(\Omega, \mathcal{H}, \mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}, \mathbb{K}^{\text{do}(\mathbb{U}, \mathbb{Q})})$ is indeed a valid causal space, satisfying the axioms of Definition 3.1.1. The following theorem does exactly that.




Figure 3.3: The calculation of the interventional causal kernels $K_S^{\text{do}(\mathbb{U}, \mathbb{Q})}(\omega; \cdot) \in \mathbb{K}^{\text{do}(\mathbb{U}, \mathbb{Q})}$ following an intervention, according to Eq. (3.2). It is the measure when the $\mathbb{S}$ -coordinates are fixed to the values ω after an intervention on $\mathcal{H}_U$. The product measure $\delta_{\omega|_{\mathbb{S}}} \otimes \mathbb{Q}|_{U \setminus \mathbb{S}}$ on the σ -algebra $\mathcal{H}_{S \cup U}$ is propagated to all of $\mathcal{H}$ via the causal kernel $K_{S \cup U}$, to give $K_S^{\text{do}(\mathbb{U}, \mathbb{Q})}(\omega; \cdot) = \int \mathbb{Q}(d\omega'_{U \setminus \mathbb{S}}) K_{S \cup U}((\omega_{\mathbb{S}}, \omega'_{U \setminus \mathbb{S}}); \cdot)$.

Theorem 3.2.3. *The interventional causal space $(\Omega, \mathcal{H}, \mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}, \mathbb{K}^{\text{do}(\mathbb{U}, \mathbb{Q})})$ is a valid causal space.*

Proof. The interventional probability measure $\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}$ is indeed a valid probability measure, by directly applying Theorem 2.5.4(i). We also prove that each $K_S^{\text{do}(\mathbb{U}, \mathbb{Q})}$ is a transition probability kernel. To this end, let us define, for each $\omega \in \Omega$, a probability measure $M(\omega; \cdot)$ on $\mathcal{H}_{S \cup U}$ as the product of δ_{ω} on $\mathcal{H}_{\mathbb{S}}$ and $\mathbb{Q}$ on $\mathcal{H}_{U \setminus \mathbb{S}}$ (see the remark after Proposition 2.4.4 for this abuse of notation). Then for any fixed cylinder $B \cap C \in \mathcal{H}_{S \cup U}$ with $B \in \mathcal{H}_{\mathbb{S}}$ and $C \in \mathcal{H}_{U \setminus \mathbb{S}}$, the map $\omega \mapsto M(\omega; B \cap C) = \mathbf{1}_B(\omega) \mathbb{Q}(C)$ is clearly $\mathcal{H}_{\mathbb{S}}$ -measurable (see Lemma 2.2.3). Hence, by Lemma 2.3.3, M is a transition probability kernel from $(\Omega, \mathcal{H}_{\mathbb{S}})$ into $(\Omega, \mathcal{H}_{S \cup U})$. But by Eq. (3.2), we have

$$\int M(\omega; d\omega') K_{S \cup U}(\omega'; A) = \int \mathbb{Q}(d\omega') K_{S \cup U}((\omega_{\mathbb{S}}, \omega'_{U \setminus \mathbb{S}}); A) = K_S^{\text{do}(\mathbb{U}, \mathbb{Q})}(\omega; A),$$

hence $K_S^{\text{do}(\mathbb{U}, \mathbb{Q})}$ is a transition probability kernel from $(\Omega, \mathcal{H}_{\mathbb{S}})$ into $(\Omega, \mathcal{H})$ by Theorem 2.5.4(iii).

We now show that $\mathbb{K}^{\text{do}(\mathbb{U}, \mathbb{Q})}$ is a causal mechanism that satisfies the two axioms of Definition 3.1.1.

(i) See that

$$K_{\emptyset}^{\text{do}(\mathbb{U}, \mathbb{Q})}(\omega; A) = \int \mathbb{Q}(d\omega') K_{\emptyset \cup \mathbb{U}}((\omega_{\emptyset}, \omega'_{\mathbb{U}}); A) = \int \mathbb{Q}(d\omega) K_{\mathbb{U}}(\omega; A) = \mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}(A)$$

by Eq. (3.1). This is precisely the trivial intervention axiom (Axiom (i)) in the interventional causal space $(\Omega, \mathcal{H}, \mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}, \mathbb{K}^{\text{do}(\mathbb{U}, \mathbb{Q})})$.

(ii) Take any $B \in \mathcal{H}_{\mathbb{S}}$. Then we can apply the interventional determinism axiom (Axiom (ii)) on the causal kernel $K_{\mathbb{S} \cup \mathbb{U}}$, since $B \in \mathcal{H}_{\mathbb{S}} \subseteq \mathcal{H}_{\mathbb{S} \cup \mathbb{U}}$, to obtain

$$K_{\mathbb{S}}^{\text{do}(\mathbb{U}, \mathbb{Q})}(\omega; B) = \int \mathbb{Q}(d\omega') K_{\mathbb{S} \cup \mathbb{U}}((\omega_{\mathbb{S}}, \omega'_{\mathbb{U} \setminus \mathbb{S}}); B) = \mathbf{1}_B(\omega) \int \mathbb{Q}(d\omega') = \mathbf{1}_B(\omega),$$

which is the interventional determinism axiom in the interventional causal space $(\Omega, \mathcal{H}, \mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}, \mathbb{K}^{\text{do}(\mathbb{U}, \mathbb{Q})})$.

□

The next result collects some properties about interventional causal kernels.

Lemma 3.2.4. *Take $\mathbb{U}, \mathbb{V} \in \mathcal{P}(\mathbb{T})$ and some assignment measures $\mathbb{Q}, \mathbb{Q}'$ on $(\Omega, \mathcal{H}_{\mathbb{U}})$ and $(\Omega, \mathcal{H}_{\mathbb{V}})$ respectively.*

(i) *Suppose that $\mathbb{U} \subseteq \mathbb{V}$. Then for all $\omega \in \Omega$ and $A \in \mathcal{H}$, we have $K_{\mathbb{V}}^{\text{do}(\mathbb{U}, \mathbb{Q})}(\omega; A) = K_{\mathbb{V}}(\omega; A)$.*

(ii) *We have $(\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})})^{\text{do}(\mathbb{V}, \mathbb{Q}')} = \mathbb{P}^{\text{do}(\mathbb{V} \cup \mathbb{U}, \mathbb{Q}' \otimes \mathbb{Q}|_{\mathbb{U} \setminus \mathbb{V}})}$ and, for all $\mathbb{S} \in \mathcal{P}(\mathbb{T})$, all $\omega \in \Omega$ and all $A \in \mathcal{H}$, $(K^{\text{do}(\mathbb{U}, \mathbb{Q})})_{\mathbb{S}}^{\text{do}(\mathbb{V}, \mathbb{Q}')}(\omega; A) = K_{\mathbb{S}}^{\text{do}(\mathbb{V} \cup \mathbb{U}, \mathbb{Q}' \otimes \mathbb{Q}|_{\mathbb{U} \setminus \mathbb{V}})}(\omega; A)$.*

(iii) *If $\mathbb{V} \cap \mathbb{U} = \emptyset$, then the interventions commute:*

$$(\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})})^{\text{do}(\mathbb{V}, \mathbb{Q}')} = (\mathbb{P}^{\text{do}(\mathbb{V}, \mathbb{Q}')})^{\text{do}(\mathbb{U}, \mathbb{Q})} \text{ and } (\mathbb{K}^{\text{do}(\mathbb{U}, \mathbb{Q})})^{\text{do}(\mathbb{V}, \mathbb{Q}')} = (\mathbb{K}^{\text{do}(\mathbb{V}, \mathbb{Q}')})^{\text{do}(\mathbb{U}, \mathbb{Q})}.$$

Proof. (i) Since $\mathbb{U} \subseteq \mathbb{V}$, we have $\mathbb{U} \setminus \mathbb{V} = \emptyset$ and $\mathbb{V} \cup \mathbb{U} = \mathbb{V}$. Hence,

$$K_{\mathbb{V}}^{\text{do}(\mathbb{U}, \mathbb{Q})}(\omega; A) = \int \mathbb{Q}(d\omega'_{\mathbb{U} \setminus \mathbb{V}}) K_{\mathbb{V} \cup \mathbb{U}}((\omega_{\mathbb{V}}, \omega'_{\mathbb{U} \setminus \mathbb{V}}); A) = K_{\mathbb{V}}(\omega; A).$$

(ii) Take any $A \in \mathcal{H}$. Then by Fubini's theorem (Theorem 2.5.6),

$$\begin{aligned} (\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})})^{\text{do}(\mathbb{V}, \mathbb{Q}')} (A) &= \int \mathbb{Q}'(d\omega) K_{\mathbb{V}}^{\text{do}(\mathbb{U}, \mathbb{Q})}(\omega; A) \\ &= \int \mathbb{Q}'(d\omega) \int \mathbb{Q}(d\omega'_{\mathbb{U} \setminus \mathbb{V}}) K_{\mathbb{V} \cup \mathbb{U}}((\omega_{\mathbb{V}}, \omega'_{\mathbb{U} \setminus \mathbb{V}}); A) \\ &= \int \mathbb{Q}' \otimes \mathbb{Q}|_{\mathbb{U} \setminus \mathbb{V}}((d\omega_{\mathbb{V}}, d\omega'_{\mathbb{U} \setminus \mathbb{V}})) K_{\mathbb{V} \cup \mathbb{U}}((\omega_{\mathbb{V}}, \omega'_{\mathbb{U} \setminus \mathbb{V}}); A) \\ &= \mathbb{P}^{\text{do}(\mathbb{V} \cup \mathbb{U}, \mathbb{Q}' \otimes \mathbb{Q}|_{\mathbb{U} \setminus \mathbb{V}})}(A). \end{aligned}$$

For the interventional causal mechanism, take any $\mathbb{S} \in \mathcal{P}(\mathbb{T})$, any $\omega \in \Omega$ and any $A \in \mathcal{H}$. Then

$$\begin{aligned}
 & (K^{\text{do}(\mathbb{U}, \mathbb{Q})}_{\mathbb{S}})^{\text{do}(\mathbb{V}, \mathbb{Q}')}(\omega; A) \\
 &= \int \mathbb{Q}'(d\omega'_{\mathbb{V} \setminus \mathbb{S}}) K_{\mathbb{S} \cup \mathbb{V}}^{\text{do}(\mathbb{U}, \mathbb{Q})}((\omega_{\mathbb{S}}, \omega'_{\mathbb{V} \setminus \mathbb{S}}); A) \\
 &= \int \mathbb{Q}'(d\omega'_{\mathbb{V} \setminus \mathbb{S}}) \int \mathbb{Q}(d\omega''_{\mathbb{U} \setminus (\mathbb{S} \cup \mathbb{V})}) K_{\mathbb{U} \cup \mathbb{S} \cup \mathbb{V}}((\omega_{\mathbb{S}}, \omega'_{\mathbb{V} \setminus \mathbb{S}}, \omega''_{\mathbb{U} \setminus (\mathbb{S} \cup \mathbb{V})}); A) \\
 &= \int \mathbb{Q}'|_{\mathbb{V} \setminus \mathbb{S}} \otimes \mathbb{Q}|_{\mathbb{U} \setminus (\mathbb{S} \cup \mathbb{V})}(d\omega'_{\mathbb{V} \setminus \mathbb{S}}, d\omega''_{\mathbb{U} \setminus (\mathbb{S} \cup \mathbb{V})}) K_{\mathbb{U} \cup \mathbb{S} \cup \mathbb{V}}((\omega_{\mathbb{S}}, \omega'_{\mathbb{V} \setminus \mathbb{S}}, \omega''_{\mathbb{U} \setminus (\mathbb{S} \cup \mathbb{V})}); A) \\
 &= K_{\mathbb{S}}^{\text{do}(\mathbb{V} \cup \mathbb{U}, \mathbb{Q}' \otimes \mathbb{Q}|_{\mathbb{U} \setminus \mathbb{V}})}(\omega; A).
 \end{aligned}$$

(iii) This is an immediate corollary of part (ii). $\square$

We give examples of interventions on homogeneous causal spaces by returning to the medicine-blood pressure example in Example 3.1.6.

Example 3.2.5. Suppose that we intervene to administer the medicine, that is to say, we intervene with the measure δ_1 on $\mathcal{H}_{\text{Med}}$. Then the probability of the post-intervention blood pressure being normal is given by

$$\mathbb{P}^{\text{do}(\text{Med}, \delta_1)}(\{\omega_{\text{Blood}} = 0\}) = K_{\text{Med}}(1; \{\omega_{\text{Blood}} = 0\}) = 0.97,$$

from the last row, third column of Table 3.2. If we want to model a randomised trial, with equal probability of giving the medicine and not giving the medicine, then we let $\mathbb{Q}(\omega_{\text{Med}} = 0) = \mathbb{Q}(\omega_{\text{Med}} = 1) = \frac{1}{2}$. Then by Eq. (3.1), the post-intervention probability of the blood pressure being normal is

$$\mathbb{P}^{\text{do}(\text{Med}, \mathbb{Q})}(\{\omega_{\text{Blood}} = 0\}) = \frac{1}{2}0.97 + \frac{1}{2}0.93 = 0.95.$$

Of course, in a randomised trial, the interest is in comparing the treated group against the control group, rather than the pooled probabilities, as is the case above. $\square$

Let us revisit Remark 3.1.7 on the product structure of a causal space.

Remark 3.2.6 (Product structure, II). (i) Taking the product structure explicitly into account is both natural and inevitable for causality, if the causal relationships are to be meaningfully decoupled from the act of intervention itself. For if events are not in separate components of a product space, one can always intervene on one of these events in such a way that the measures on the other events have to change.

Considering the measurable spaces in Remark 3.1.7, if we intervene on $(\Omega_3, \mathcal{E}_3)$ by giving measure 1 to the event $\{1\}$, then the event $\{2, 3, 4\}$ has no choice but to have measure zero. In contrast, in the first causal space, whatever assignment measure $\mathbb{Q}$ we intervene with on $(\Omega, \mathcal{H}_1)$, the interventional measure $\mathbb{P}^{\text{do}(1, \mathbb{Q})}$ on $(\Omega, \mathcal{H}_2)$ can be any valid probability measure depending on the causal kernel K_1 . Thus, in the former case, the act of intervention and the causal effect on the rest of the space are not decoupled, whereas in the latter case, they are.

- (ii) Interventions are only defined on σ -algebras of the form $\mathcal{H}_S$ for some $S \in \mathcal{P}(\mathbb{T})$, i.e. some coordinate set of the product structure. This stands in contrast with *conditioning*, which can be done on any sub- σ -algebra $\mathcal{G} \subseteq \mathcal{H}$, not necessarily of the form $\mathcal{H}_S$ for some $S \in \mathcal{P}(\mathbb{T})$. This means that $\mathcal{G} \subseteq \mathcal{H}$ can be a sub- σ -algebra of some component $\mathcal{H}_S$ without coinciding with it. Throughout the monograph, we will clearly state whether a σ -algebra under consideration is a component of the product (denoted as $\mathcal{H}_S$, $\mathcal{H}_U$, etc. for $S, U \in \mathcal{P}(\mathbb{T})$), or an arbitrary sub- σ -algebra of $\mathcal{H}$ (denoted as $\mathcal{F}$, $\mathcal{G} \subseteq \mathcal{H}$ etc.).

The next result has a simple proof, that is almost tautological from Definition 3.2.1, but captures arguably the most fundamental idea in causality—that, when an intervention (experiment, controlled trial) can be conducted, then the corresponding causation can be inferred directly from the resulting interventional measure. The following result formalises this in the case of homogeneous interventions, and we call it, rather grandiosely, the *fundamental lemma of causality*.

Lemma 3.2.7 (Fundamental lemma of causality). *Suppose that an intervention takes place on $\mathcal{H}_U$ via an assignment measure $\mathbb{Q}$. Then for any events $A \in \mathcal{H}$ and $B \in \mathcal{H}_U$, we have*

$$\mathbb{E}^{\text{do}(\mathbb{U}, \mathbb{Q})}[\mathbf{1}_B(\cdot)K_U(\cdot; A)] = \mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}(A \cap B);$$

in other words, the map $\omega \mapsto K_U(\omega; A)$ is a version of the conditional interventional probability measure $\mathbb{P}_{\mathcal{H}_U}^{\text{do}(\mathbb{U}, \mathbb{Q})}(\cdot, A)$.

The idea is that we are interested in the quantity $K_U(\omega; A)$, which is the measure on the event A after intervening on the U -coordinates by fixing it to the value ω . This is a causal quantity which, a priori, has no relationship with the observational measure $\mathbb{P}$, and therefore nothing about it can be inferred from $\mathbb{P}$. Lemma 3.2.7 says that, if we are able to perform an intervention on $\mathcal{H}_U$, then we do not need to assume any relationship between $\mathbb{P}$ and K_U , because the resulting interventional measure $\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}$ tells us about K_U . This is the core idea behind the mantra that “*randomised control trials are the gold standard in causal inference*”. The lemma holds for *any* assignment measure $\mathbb{Q}$, but of course, we can only know $K_U(\cdot; A)$ on events of positive $\mathbb{Q}$ -measure.

Proof. Take any events $A \in \mathcal{H}$ and $B \in \mathcal{H}_U$. First see that $\omega \mapsto \mathbf{1}_B(\omega)K_U(\omega; A)$ is $\mathcal{H}_U$ -measurable by Lemma 2.2.7 as a product of two $\mathcal{H}_U$ -measurable functions, so we can apply Lemma 3.2.2 to see that

$$\int \mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}(d\omega) \mathbf{1}_B(\omega)K_U(\omega; A) = \int \mathbb{Q}(d\omega) \mathbf{1}_B(\omega)K_U(\omega; A) = \int \mathbb{Q}(d\omega)K_U(\omega; A \cap B),$$

where the last equality is obtained by the interventional determinism axiom (Theorem 3.1.2(a)). The last expression is the definition of $\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}(A \cap B)$. $\square$

3.3 Causal effects

It is often said that the notion of *independence* is what truly sets probability theory apart from measure theory. In terms of axioms, the only additional requirement that

probability measures have over general measures is that they be normalised to one. Beyond the two requirements of Axioms (i) and (ii), the causal kernels are also simply transition probability kernels, which already exist in probability theory. So what are the notions in causal space theory that truly set the theory of causality apart from probability theory, in the same way that independence gets probability theory started?

In this section, we introduce the first of the two such fundamental notions, namely, what it means for a causal kernel to have a *causal effect* on an event. In the next section, we will introduce the second notion, called *sources*. These two notions link the causal kernels to the observational measure $\mathbb{P}$ or to each other; the causal kernels, as discussed in Section 3.1, are primitive objects that are a priori not related to $\mathbb{P}$ or to each other in any way.

Definition 3.3.1 (Causal effect). There are three mutually exclusive cases of causal effects. Let us take $\mathbb{U} \in \mathcal{P}(\mathbb{T})$, an event $A \in \mathcal{H}$ and a sub- σ -algebra $\mathcal{F} \subseteq \mathcal{H}$.

- We say that $K_{\mathbb{U}}$ has an *active causal effect* on $A \in \mathcal{H}$ if there exists some $\omega \in \Omega$ such that $K_{\mathbb{U}}(\omega; A) \neq \mathbb{P}(A)$.

We say that $K_{\mathbb{U}}$ has an active causal effect on $\mathcal{F}$ if $K_{\mathbb{U}}$ has an active causal effect on some $A \in \mathcal{F}$.

- We say that the causal kernel $K_{\mathbb{U}}$ has *no causal effect* on A if, for all $\mathbb{S} \in \mathcal{P}(\mathbb{T})$, we have $K_{\mathbb{S}}(\omega; A) = K_{\mathbb{S} \setminus \mathbb{U}}(\omega; A)$ for all $\omega \in \Omega$.

We say that $K_{\mathbb{U}}$ has no causal effect on $\mathcal{F}$ if $K_{\mathbb{U}}$ has no causal effect on every $A \in \mathcal{F}$.

- Otherwise, we say that $K_{\mathbb{U}}$ has a *dormant causal effect* on A (resp. $\mathcal{F}$).

If $K_{\mathbb{U}}$ has either an active or a dormant causal effect on A , then we sometimes say that $K_{\mathbb{U}}$ has a *causal effect* on A .

Let us walk through each notion in order. The active causal effect is the simplest to understand; it simply asks whether there is some $\omega \in \Omega$ such that, by fixing the $\mathbb{U}$ -coordinates to ω , the probability measure of A changes from the observational measure $\mathbb{P}$. If $K_{\mathbb{U}}$ has an active causal effect on A , we can see that, by simply intervening with the assignment measure δ_{ω} , where ω is the outcome for which $K_{\mathbb{U}}(\omega; A) \neq \mathbb{P}(A)$, we have an interventional measure that is different from the observational:

$$\mathbb{P}^{\text{do}(\mathbb{U}, \delta_{\omega})}(A) = \int \delta_{\omega}(d\omega') K_{\mathbb{U}}(\omega'; A) = K_{\mathbb{U}}(\omega; A) \neq \mathbb{P}(A).$$

On the other hand, if $K_{\mathbb{U}}$ does not have an active causal effect on A , then for any assignment measure $\mathbb{Q}$ on $(\Omega, \mathcal{H}_{\mathbb{U}})$, we have

$$\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}(A) = \int \mathbb{Q}(d\omega) K_{\mathbb{U}}(\omega; A) = \int \mathbb{Q}(d\omega) \mathbb{P}(A) = \mathbb{P}(A); \quad (3.3)$$

in other words, upon intervention on the $\mathbb{U}$ -coordinates, the measure on the event A does not change from the observational measure. The absence of an active causal effect of $K_{\mathbb{U}}$




Figure 3.4: Active causal effect (Definition 3.3.1). If the probability of A remains at $\mathbb{P}(A)$ upon fixing the $\mathbb{U}$ -coordinates to any value $\omega_{\mathbb{U}}$, then $K_{\mathbb{U}}$ does not have an active causal effect on A .

on A is sufficient to ensure $\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}(A) = \mathbb{P}(A)$. However, it is not a *necessary* condition. It may be that there are outcomes $\omega \in \Omega$ for which we have $K_{\mathbb{U}}(\omega; A) \neq \mathbb{P}(A)$, but $\mathbb{Q}$ is such that those differences are averaged out to yield $\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}(A) = \mathbb{P}(A)$. The definition rules out such accidental coincidence of $\mathbb{P}(A)$ and $\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}(A)$, and only declares an absence of an active causal effect if $K_{\mathbb{U}}(\omega; A) = \mathbb{P}(A)$ *for all* $\omega \in \Omega$.

The notion of no causal effect insists on more than for $K_{\mathbb{U}}$ not to have an active causal effect. While active causal effect only considers interventions on the $\mathbb{U}$ -coordinates, no effect considers interventions on *all* possible sets of coordinates $\mathbb{S}$, and only declares that $K_{\mathbb{U}}$ has no causal effect on A if the $\mathbb{U}$ -coordinates never make a difference for an intervention on any $\mathbb{S}$. Note that, if $K_{\mathbb{U}}$ has no causal effect on A , then no $K_{\mathbb{S}}$ with $\mathbb{S} \subseteq \mathbb{U}$ can have an active causal effect on A : by applying the trivial intervention axiom (Axiom (i)):

$$K_{\mathbb{S}}(\omega; A) = K_{\mathbb{S} \setminus \mathbb{U}}(\omega; A) = K_{\emptyset}(\omega; A) = \mathbb{P}(A).$$

In particular, if $K_{\mathbb{U}}$ has no causal effect on A , it implies that $K_{\mathbb{U}}$ has no active causal effect on A .

Now let us have a look at what the axioms imply for causal effects.

- (i) The trivial intervention axiom (Axiom (i)) gives $K_{\emptyset}(\omega; A) = \mathbb{P}(A)$ for all $\omega \in \Omega$ and all $A \in \mathcal{H}$, so $K_{\emptyset}$ does not have an active causal effect on any $A \in \mathcal{H}$. In fact, for all $\mathbb{S} \in \mathcal{P}(\mathbb{T})$ and all $\omega \in \Omega$, we have $K_{\mathbb{S}}(\omega; A) = K_{\mathbb{S} \setminus \emptyset}(\omega; A)$, hence $K_{\emptyset}$ has no causal effect on any $A \in \mathcal{H}$. The latter is not reliant on the trivial intervention axiom.
- (ii) The interventional determinism axiom (Axiom (ii)) implies, provided $\mathcal{H}_{\mathbb{U}}$ is non-trivial, that the kernel $K_{\mathbb{U}}$ has an active causal effect on $\mathcal{H}_{\mathbb{U}}$. Indeed, take any $B \in \mathcal{H}_{\mathbb{U}}$ with $B \neq \emptyset$ and $B \neq \Omega$. Then if $0 < \mathbb{P}(B) < 1$, we have $K_{\mathbb{U}}(\omega; B) =$




Figure 3.5: No causal effect (Definition 3.3.1). If fixing any $\mathbb{S}$ -coordinates to any value $\omega_{\mathbb{S}}$ is the same as fixing the $(\mathbb{S}\setminus\mathbb{U})$ -coordinates to the same (marginal) values $\omega_{\mathbb{S}\setminus\mathbb{U}}$ (i.e. $K_{\mathbb{U}}(\omega; A) = K_{\mathbb{S}\setminus\mathbb{U}}(\omega; A)$ for all $\mathbb{S} \in \mathcal{P}(\mathbb{T})$ and all $\omega \in \Omega$), then $K_{\mathbb{U}}$ has no causal effect on A .

$\mathbf{1}_B(\omega) \neq \mathbb{P}(B)$ for any $\omega \in \Omega$; if $\mathbb{P}(B) = 1$, then $K_{\mathbb{U}}(\omega; B) = 0 \neq \mathbb{P}(B)$ for $\omega \notin B$; and if $\mathbb{P}(B) = 0$, then $K_{\mathbb{U}}(\omega; B) = 1 \neq \mathbb{P}(B)$ for $\omega \in B$.

Note also that, unlike (in)dependence, causal effect is an asymmetric notion. If $\mathcal{H}_{\mathbb{U}}$ is independent of $\mathcal{H}_{\mathbb{V}}$, then $\mathcal{H}_{\mathbb{V}}$ is independent of $\mathcal{H}_{\mathbb{U}}$. However, it is of course possible that $K_{\mathbb{U}}$ has a causal effect on $\mathcal{H}_{\mathbb{V}}$, while $K_{\mathbb{V}}$ has no causal effect on $\mathcal{H}_{\mathbb{U}}$. It is also possible for a $K_{\mathbb{U}}$ not to have an active causal effect on any event in $\mathcal{H}_{\mathbb{T}\setminus\mathbb{U}}$. The following example demonstrates both of these cases with active causal effects.

Example 3.3.2. Take some $\mathbb{U} \in \mathcal{P}(\mathbb{T})$. Then $K_{\mathbb{U}}(\omega; \cdot) = (\delta_{\omega}|_{\mathbb{U}} \otimes \mathbb{P}|_{\mathbb{T}\setminus\mathbb{U}})(\cdot)$ is a valid causal kernel. Indeed, if $\mathbb{U} = \emptyset$, then $K_{\mathbb{U}}(\omega; A) = \mathbb{P}(A)$, so Axiom (i) is satisfied, and for general $\mathbb{U} \in \mathcal{P}(\mathbb{T})$, for any $B \in \mathcal{H}_{\mathbb{U}}$, we have $K_{\mathbb{U}}(\omega; B) = \delta_{\omega}(B)$, so Axiom (ii) is satisfied.

This causal kernel does not have an active causal effect on any $A \in \mathcal{H}_{\mathbb{T}\setminus\mathbb{U}}$: for any $\omega \in \Omega$, we have $K_{\mathbb{U}}(\omega; A) = \mathbb{P}(A)$. This construction works for *any* $\mathbb{P}$, so it is worth remarking explicitly that any observational measure admits a causation-free kernel.

Further, if $K_{\mathbb{T}\setminus\mathbb{U}}$ is such that $K_{\mathbb{T}\setminus\mathbb{U}}(\omega; B) \neq \mathbb{P}(B)$ for some $\omega \in \Omega$ and $B \in \mathcal{H}_{\mathbb{U}}$, then $K_{\mathbb{T}\setminus\mathbb{U}}$ has an active causal effect on $\mathcal{H}_{\mathbb{U}}$. $\square$

Dormant causal effect is the remaining case, whereby the probability measure on A does not change under any intervention on the $\mathbb{U}$ -coordinates, but there is some $\mathbb{S} \in \mathcal{P}(\mathbb{T})$ such that intervening on the $\mathbb{S}$ -coordinates depends on the $\mathbb{U}$ -coordinates. In other words, (a subset of) the $\mathbb{U}$ -coordinates become active only under interventions on a different set of coordinates $\mathbb{S}$. The following example presents a toy case of this phenomenon.

Example 3.3.3. Let us take three binary outcome sets $\Omega_1 = \Omega_2 = \Omega_3 = \{0, 1\}$, so that the outcome set $\Omega = \Omega_1 \times \Omega_2 \times \Omega_3$ has 8 elements. Let $\mathbb{P}$ be the uniform observational measure, and let us specify a subset of the causal kernels as in Table 3.3. The marginal

Outcome	$\omega = (\omega_1, \omega_2, \omega_3)$							
	(0, 0, 0)	(1, 0, 0)	(0, 1, 0)	(0, 0, 1)	(1, 1, 0)	(1, 0, 1)	(0, 1, 1)	(1, 1, 1)
$\mathbb{P}$	1/8	1/8	1/8	1/8	1/8	1/8	1/8	1/8
$K_{\{1\}}(0; \cdot)$	1/4	0	1/4	1/4	0	0	1/4	0
$K_{\{1\}}(1; \cdot)$	0	1/4	0	0	1/4	1/4	0	1/4
$K_{\{2\}}(0; \cdot)$	1/8	1/8	0	3/8	0	3/8	0	0
$K_{\{1,2\}}((0, 0); \cdot)$	1/8	0	0	7/8	0	0	0	0

Table 3.3: In Example 3.3.3, $K_{\{1\}}$ has a dormant causal effect on $\mathcal{H}_3$ and $K_{\{2\}}$ has an active causal effect on $\mathcal{H}_3$.

observational measure on Ω_3 (first row of Table 3.3) is

$$\mathbb{P}(\omega_3 = 0) = \mathbb{P}(\omega_3 = 1) = 1/2.$$

Intervening on $\mathcal{H}_1$ with $\omega_1 = 0$ or $\omega_1 = 1$ keeps the measure on $\mathcal{H}_2$ and $\mathcal{H}_3$ uniform (second and third rows of Table 3.3):

$$K_1(0; \{\omega_3 = 0\}) = K_1(0; \{\omega_3 = 1\}) = 1/2 = K_1(1; \{\omega_3 = 0\}) = K_1(1; \{\omega_3 = 1\}),$$

and so K_1 has no active causal effect on $\mathcal{H}_2$ and $\mathcal{H}_3$. On the other hand, we can see that K_2 has an active causal effect on $\mathcal{H}_3$, since

$$K_2(0; \{\omega_3 = 0\}) = 1/4 \neq 1/2 = \mathbb{P}(\omega_3 = 0).$$

Finally, $K_{1,2}$ has an active causal effect on $\mathcal{H}_3$, since

$$K_{1,2}((0, 0); \{\omega_3 = 0\}) = 1/8 \neq 1/2 = \mathbb{P}(\omega_3 = 0),$$

and in particular, as $K_{1,2}((0, 0); \{\omega_3 = 0\}) = 1/8 \neq 1/4 = K_2(0; \{\omega_3 = 0\})$, we see that K_1 has a dormant causal effect on $\{\omega_3 = 0\}$: the intervention $\omega_1 = 0$ has no active causal effect by itself, but the joint intervention $\omega_{1,2} = (0, 0)$ is different to the intervention $\omega_2 = 0$. For K_1 to have no causal effect on $\mathcal{H}_3$, we would, in particular, have required $K_{1,2}((\omega_1, \omega_2); A) = K_2(\omega_2; A)$ for all $\omega_1 \in \Omega_1$, all $\omega_2 \in \Omega_2$ and all $A \in \mathcal{H}_3$. $\square$

Here are some results that are easy to deduce.

Lemma 3.3.4. *Suppose that $\mathbb{U}, \mathbb{V} \in \mathcal{P}(\mathbb{T})$, and take an event $A \in \mathcal{H}$.*

- (i) *If $\mathbb{V} \subseteq \mathbb{U}$ and $K_{\mathbb{U}}$ has no causal effect on A , then $K_{\mathbb{V}}$ has no causal effect on A .*
- (ii) *If neither $K_{\mathbb{U}}$ nor $K_{\mathbb{V}}$ has a causal effect on A , then $K_{\mathbb{U} \cup \mathbb{V}}$ has no causal effect on A .*
- (iii) *For any assignment measure $\mathbb{Q}$ on $(\Omega, \mathcal{H}_{\mathbb{V}})$ and for any $\mathbb{U} \in \mathcal{P}(\mathbb{T})$ with $\mathbb{U} \cap \mathbb{V} = \emptyset$, $K_{\mathbb{U}}^{\text{do}(\mathbb{V}, \mathbb{Q})}$ has no causal effect on $\mathcal{H}_{\mathbb{V}}$.*
- (iv) *If $K_{\mathbb{U}}$ has no causal effect on A , then for any $\mathbb{V} \in \mathcal{P}(\mathbb{T})$ and any assignment measure $\mathbb{Q}$ on $(\Omega, \mathcal{H}_{\mathbb{V}})$, the interventional causal kernel $K_{\mathbb{U}}^{\text{do}(\mathbb{V}, \mathbb{Q})}$ has no causal effect on A .*

Proof. (i) See that, for any $\mathbb{S} \in \mathcal{P}(\mathbb{T})$ and any $\omega \in \Omega$,

$$K_{\mathbb{S} \setminus \mathbb{V}}(\omega; A) = K_{(\mathbb{S} \setminus \mathbb{V}) \setminus \mathbb{U}}(\omega; A) = K_{\mathbb{S} \setminus \mathbb{U}}(\omega; A) = K_{\mathbb{S}}(\omega; A),$$

where we twice used the fact that $K_{\mathbb{U}}$ has no causal effect on A , first with $\mathbb{S} \setminus \mathbb{V}$ and then with $\mathbb{S}$.

(ii) See that, for any $\mathbb{S} \in \mathcal{P}(\mathbb{T})$ and any $\omega \in \Omega$,

$$K_{\mathbb{S} \setminus (\mathbb{U} \cup \mathbb{V})}(\omega; A) = K_{(\mathbb{S} \setminus \mathbb{U}) \setminus \mathbb{V}}(\omega; A) = K_{\mathbb{S} \setminus \mathbb{U}}(\omega; A) = K_{\mathbb{S}}(\omega; A),$$

where we first used the fact that $K_{\mathbb{V}}$ has no causal effect on A with $\mathbb{S} \setminus \mathbb{U}$, then that $K_{\mathbb{U}}$ has no causal effect on A with $\mathbb{S}$.

(iii) Take any $\mathbb{S} \in \mathcal{P}(\mathbb{T})$ and any $B \in \mathcal{H}_{\mathbb{V}}$. Then recalling the definition of interventional causal kernels from Eq. (3.2),

$$\begin{aligned} K_{\mathbb{S}}^{\text{do}(\mathbb{V}, \mathbb{Q})}(\omega; B) &= \int \mathbb{Q}(d\omega'_{\mathbb{V} \setminus \mathbb{S}}) K_{\mathbb{S} \cup \mathbb{V}}((\omega_{\mathbb{S}}, \omega'_{\mathbb{V} \setminus \mathbb{S}}); B) \\ &= \int \mathbb{Q}(d\omega'_{\mathbb{V} \setminus \mathbb{S}}) \mathbf{1}_B((\omega_{\mathbb{S} \cap \mathbb{V}}, \omega'_{\mathbb{V} \setminus \mathbb{S}})), \end{aligned}$$

by the interventional determinism axiom, since $B \in \mathcal{H}_{\mathbb{V}} \subseteq \mathcal{H}_{\mathbb{S} \cup \mathbb{V}}$. Since $\mathbb{U} \cap \mathbb{V} = \emptyset$, we have

$$\mathbb{V} \setminus \mathbb{S} = \mathbb{V} \setminus (\mathbb{S} \setminus \mathbb{U}) \quad \text{and} \quad \mathbb{S} \cap \mathbb{V} = (\mathbb{S} \setminus \mathbb{U}) \cap \mathbb{V}.$$

Hence,

$$\begin{aligned} K_{\mathbb{S}}^{\text{do}(\mathbb{V}, \mathbb{Q})}(\omega; B) &= \int \mathbb{Q}(d\omega'_{\mathbb{V} \setminus (\mathbb{S} \setminus \mathbb{U})}) \mathbf{1}_B((\omega_{(\mathbb{S} \setminus \mathbb{U}) \cap \mathbb{V}}, \omega'_{\mathbb{V} \setminus (\mathbb{S} \setminus \mathbb{U})})) \\ &= \int \mathbb{Q}(d\omega'_{\mathbb{V} \setminus (\mathbb{S} \setminus \mathbb{U})}) K_{(\mathbb{S} \setminus \mathbb{U}) \cup \mathbb{V}}((\omega_{\mathbb{S} \setminus \mathbb{U}}, \omega'_{\mathbb{V} \setminus (\mathbb{S} \setminus \mathbb{U})}); B) \\ &= K_{\mathbb{S} \setminus \mathbb{U}}^{\text{do}(\mathbb{V}, \mathbb{Q})}(\omega; B), \end{aligned}$$

where we applied the interventional determinism axiom again on the second line, since $B \in \mathcal{H}_{\mathbb{V}} \subseteq \mathcal{H}_{(\mathbb{S} \setminus \mathbb{U}) \cup \mathbb{V}}$.

(iv) Take any $\mathbb{S} \in \mathcal{P}(\mathbb{T})$ and any $\omega \in \Omega$. Then

$$\begin{aligned} K_{\mathbb{S}}^{\text{do}(\mathbb{V}, \mathbb{Q})}(\omega; A) &= \int \mathbb{Q}(d\omega'_{\mathbb{V} \setminus \mathbb{S}}) K_{\mathbb{S} \cup \mathbb{V}}((\omega_{\mathbb{S}}, \omega'_{\mathbb{V} \setminus \mathbb{S}}); A) \\ &= \int \mathbb{Q}(d\omega'_{(\mathbb{V} \setminus \mathbb{S}) \setminus \mathbb{U}}) K_{(\mathbb{S} \cup \mathbb{V}) \setminus \mathbb{U}}((\omega_{\mathbb{S} \setminus \mathbb{U}}, \omega'_{(\mathbb{V} \setminus \mathbb{S}) \setminus \mathbb{U}}); A), \end{aligned}$$

where we used the fact that $K_{\mathbb{U}}$ has no causal effect on A with $\mathbb{S} \cup \mathbb{V}$. But here, we have

$$(\mathbb{S} \cup \mathbb{V}) \setminus \mathbb{U} = ((\mathbb{S} \setminus \mathbb{U}) \cup \mathbb{V}) \setminus \mathbb{U}, \quad \text{and} \quad (\mathbb{V} \setminus \mathbb{S}) \setminus \mathbb{U} = (\mathbb{V} \setminus (\mathbb{S} \setminus \mathbb{U})) \setminus \mathbb{U},$$

and so

$$\begin{aligned} K_S^{\text{do}(\mathbb{V}, \mathbb{Q})}(\omega; A) &= \int \mathbb{Q}(d\omega'_{(\mathbb{V} \setminus (\mathbb{S} \setminus \mathbb{U})) \setminus \mathbb{U}}) K_{((\mathbb{S} \setminus \mathbb{U}) \cup \mathbb{V}) \setminus \mathbb{U}}(\omega_{\mathbb{S} \setminus \mathbb{U}}, \omega'_{(\mathbb{V} \setminus (\mathbb{S} \setminus \mathbb{U})) \setminus \mathbb{U}}; A) \\ &= \int \mathbb{Q}(d\omega'_{\mathbb{V} \setminus (\mathbb{S} \setminus \mathbb{U})}) K_{(\mathbb{S} \setminus \mathbb{U}) \cup \mathbb{V}}(\omega_{\mathbb{S} \setminus \mathbb{U}}, \omega'_{\mathbb{V} \setminus (\mathbb{S} \setminus \mathbb{U})}; A) \\ &= K_{\mathbb{S} \setminus \mathbb{U}}^{\text{do}(\mathbb{V}, \mathbb{Q})}(\omega; A), \end{aligned}$$

where, on the second line, we used the fact that $K_{\mathbb{U}}$ has no causal effect on A with $(\mathbb{S} \setminus \mathbb{U}) \cup \mathbb{V}$. This is precisely the definition that $K_{\mathbb{U}}^{\text{do}(\mathbb{V}, \mathbb{Q})}$ has no causal effect on A . $\square$

The contrapositive of part (i) means that, if $K_{\mathbb{V}}$ has a causal effect on A and $\mathbb{V} \subseteq \mathbb{U}$, then $K_{\mathbb{U}}$ must have a causal effect on A . But part (i) cannot be applied to active causal effects. Indeed, even if $\mathbb{V} \subseteq \mathbb{U}$ and $K_{\mathbb{U}}$ has no active causal effect on A , it does not mean that $K_{\mathbb{V}}$ cannot have an active causal effect on A . The contrapositive of part (ii) means that, if $K_{\mathbb{U} \cup \mathbb{V}}$ has a causal effect on A , then either $K_{\mathbb{U}}$ or $K_{\mathbb{V}}$ has a causal effect on A .

In the following, we give a few equivalent statements for the absence of an active causal effect, and no causal effect. In particular, the absence of an active causal effect can be characterised by the probabilistic notion of *independence* in the interventional causal space.

Theorem 3.3.5. *Let $\mathbb{U} \in \mathcal{P}(\mathbb{T})$ and $A \in \mathcal{H}$.*

(i) *The following statements are equivalent.*

- (a) $K_{\mathbb{U}}$ does not have an active causal effect on A .
- (b) For all $\omega \in \Omega$ and all $B \in \mathcal{H}_{\mathbb{U}}$, we have $K_{\mathbb{U}}(\omega; A \cap B) = \mathbb{P}(A) \mathbf{1}_B(\omega)$.
- (c) For any assignment measure $\mathbb{Q}$ on $(\Omega, \mathcal{H}_{\mathbb{U}})$ and any event $B \in \mathcal{H}_{\mathbb{U}}$, we have $\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}(A \cap B) = \mathbb{P}(A) \mathbb{Q}(B)$.
- (d) For any assignment measure $\mathbb{Q}$ on $(\Omega, \mathcal{H}_{\mathbb{U}})$, we have $\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}(A) = \mathbb{P}(A)$ and under the interventional measure $\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}$, the event A and the σ -algebra $\mathcal{H}_{\mathbb{U}}$ are independent.

(ii) *The following statements are equivalent.*

- (a) $K_{\mathbb{U}}$ has no causal effect on A .
- (b) For all $\mathbb{S} \in \mathcal{P}(\mathbb{T})$, all $\omega \in \Omega$ and all $B \in \mathcal{H}_{\mathbb{S} \cap \mathbb{U}}$, we have $K_{\mathbb{S}}(\omega; A \cap B) = \mathbf{1}_B(\omega) K_{\mathbb{S} \setminus \mathbb{U}}(\omega; A)$.
- (c) For any $\mathbb{S} \in \mathcal{P}(\mathbb{T})$, any measure $\mathbb{Q}$ on $(\Omega, \mathcal{H}_{\mathbb{S} \setminus \mathbb{U}})$, any measure $\mathbb{Q}'$ on $(\Omega, \mathcal{H}_{\mathbb{S} \cap \mathbb{U}})$ and any $B \in \mathcal{H}_{\mathbb{S} \cap \mathbb{U}}$, we have $\mathbb{P}^{\text{do}(\mathbb{S}, \mathbb{Q} \otimes \mathbb{Q}')} (A \cap B) = \mathbb{P}^{\text{do}(\mathbb{S} \setminus \mathbb{U}, \mathbb{Q})}(A) \mathbb{P}^{\text{do}(\mathbb{S} \cap \mathbb{U}, \mathbb{Q}')} (B)$.

Proof. (i) **(a) $\implies$ (b)** Take any $\omega \in \Omega$ and any $B \in \mathcal{H}_{\mathbb{U}}$. Then applying the interventional determinism axiom (Theorem 3.1.2(a)), it is immediate that

$$K_{\mathbb{U}}(\omega; A \cap B) = \mathbf{1}_B(\omega) K_{\mathbb{U}}(\omega; A) = \mathbf{1}_B(\omega) \mathbb{P}(A).$$

- (b) $\implies$ (c)** Take any assignment measure $\mathbb{Q}$ on $(\Omega, \mathcal{H}_U)$ and any $B \in \mathcal{H}_U$. Then by (b),

$$\begin{aligned}\mathbb{P}^{\text{do}(U, \mathbb{Q})}(A \cap B) &= \int \mathbb{Q}(d\omega) K_U(\omega; A \cap B) \\ &= \int \mathbb{Q}(d\omega) \mathbf{1}_B(\omega) \mathbb{P}(A) = \mathbb{Q}(B) \mathbb{P}(A).\end{aligned}$$

- (c) $\implies$ (d)** For the first statement, simply let $B = \Omega$ in (c). For the second statement, take any $B \in \mathcal{H}_U$. Then by (c) again,

$$\mathbb{P}^{\text{do}(U, \mathbb{Q})}(A \cap B) = \mathbb{P}(A) \mathbb{Q}(B) = \mathbb{P}^{\text{do}(U, \mathbb{Q})}(A) \mathbb{P}^{\text{do}(U, \mathbb{Q})}(B),$$

where we know that $\mathbb{P}^{\text{do}(U, \mathbb{Q})}(B) = \mathbb{Q}(B)$ from Section 3.2. Since $B \in \mathcal{H}_U$ was arbitrary, A and $\mathcal{H}_U$ are independent under $\mathbb{P}^{\text{do}(U, \mathbb{Q})}$.

- (d) $\implies$ (a)** Suppose for contradiction that K_U has an active causal effect on A . Then there exists some $\omega \in \Omega$ such that $\mathbb{P}^{\text{do}(U, \delta_\omega)}(A) = K_U(\omega; A) \neq \mathbb{P}(A)$, which directly contradicts the first statement of (d).
- (ii) **(a) $\implies$ (b)** Take any $\mathbb{S} \in \mathcal{P}(\mathbb{T})$, any $\omega \in \Omega$ and any $B \in \mathcal{H}_{\mathbb{S} \cap U}$. Then we can apply the interventional determinism axiom (Theorem 3.1.2(a)), since $B \in \mathcal{H}_{\mathbb{S} \cap U} \subseteq \mathcal{H}_{\mathbb{S}}$, to get

$$K_{\mathbb{S}}(\omega; A \cap B) = \mathbf{1}_B(\omega) K_{\mathbb{S}}(\omega; A) = \mathbf{1}_B(\omega) K_{\mathbb{S} \setminus U}(\omega; A).$$

- (b) $\implies$ (c)** Take any $\mathbb{S} \in \mathcal{P}(\mathbb{T})$, any $\mathbb{Q}$ on $(\Omega, \mathcal{H}_{\mathbb{S} \setminus U})$, any $\mathbb{Q}'$ on $(\Omega, \mathcal{H}_{\mathbb{S} \cap U})$ and any $B \in \mathcal{H}_{\mathbb{S} \cap U}$. Then by applying (b), the interventional determinism axiom and Fubini's theorem (Theorem 2.5.6),

$$\begin{aligned}\mathbb{P}^{\text{do}(\mathbb{S}, \mathbb{Q} \otimes \mathbb{Q}')} (A \cap B) &= \int \mathbb{Q} \otimes \mathbb{Q}'(d\omega) K_{\mathbb{S}}(\omega; A \cap B) \\ &= \int \mathbb{Q} \otimes \mathbb{Q}'((d\omega_{\mathbb{S} \cap U}, \omega'_{\mathbb{S} \setminus U})) \mathbf{1}_B(\omega_{\mathbb{S} \cap U}) K_{\mathbb{S} \setminus U}(\omega'_{\mathbb{S} \setminus U}; A) \\ &= \int \mathbb{Q}(d\omega_{\mathbb{S} \setminus U}) K_{\mathbb{S} \setminus U}(\omega_{\mathbb{S} \setminus U}; A) \int \mathbb{Q}'(d\omega') K_{\mathbb{S} \cap U}(\omega'_{\mathbb{S} \cap U}; B) \\ &= \mathbb{P}^{\text{do}(\mathbb{S} \setminus U, \mathbb{Q})}(A) \mathbb{P}^{\text{do}(\mathbb{S} \cap U, \mathbb{Q}')} (B).\end{aligned}$$

- (c) $\implies$ (a)** Take any $\mathbb{S} \in \mathcal{P}(\mathbb{T})$ and any $\omega \in \Omega$. Let $B = \Omega$ and $\mathbb{Q} = \delta_\omega$ on $(\Omega, \mathcal{H}_{\mathbb{S} \setminus U})$. Then by applying (c) with $\mathbb{Q}' = \delta_\omega$ on $(\Omega, \mathcal{H}_{\mathbb{S} \cap U})$, we have

$$K_{\mathbb{S}}(\omega; A) = \mathbb{P}^{\text{do}(\mathbb{S}, \delta_\omega \otimes \delta_\omega)}(A) = \mathbb{P}^{\text{do}(\mathbb{S} \setminus U, \delta_\omega)}(A) = K_{\mathbb{S} \setminus U}(\omega; A).$$

□

Theorem 3.3.5(i) is a satisfying characterisation of the absence of an active causal effect. It says: after intervention, the $\mathbb{U}$ -coordinates contain no information about A , and the marginal probability of A has not changed from the observational state. The following corollary is immediate.

Corollary 3.3.6. *The causal kernel $K_{\mathbb{U}}$ does not have an active causal effect on a σ -algebra $\mathcal{F} \subseteq \mathcal{H}$ if and only if, for any assignment measure $\mathbb{Q}$ on $(\Omega, \mathcal{H}_{\mathbb{U}})$, the measures $\mathbb{P}$ and $\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}$ coincide on $\mathcal{F}$ and the σ -algebras $\mathcal{F}$ and $\mathcal{H}_{\mathbb{U}}$ are independent under the interventional measure $\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}$.*

Finally, we show that the collection of events for which a $K_{\mathbb{U}}$ has no (active) causal effect forms a λ -system. This will be useful when checking that $K_{\mathbb{U}}$ has no (active) causal effect on a σ -algebra $\mathcal{F}$: instead of checking every event in that $\mathcal{F}$, the monotone class theorem (Theorem 2.1.5) allows us to only check events in a generating π -system.

Lemma 3.3.7. *For any $\mathbb{U} \in \mathcal{P}(\mathbb{T})$, the collections*

$$\begin{aligned}\mathcal{C} &= \{A \in \mathcal{H} : K_{\mathbb{U}}(\omega; A) = \mathbb{P}(A) \text{ for all } \omega \in \Omega\}, \\ \mathcal{D} &= \{A \in \mathcal{H} : K_{\mathbb{S}}(\omega; A) = K_{\mathbb{S} \setminus \mathbb{U}}(\omega; A) \text{ for all } \mathbb{S} \in \mathcal{P}(\mathbb{T}) \text{ and } \omega \in \Omega\}\end{aligned}$$

are λ -systems.

Proof. We go through the defining properties of a λ -system in order.

- For all $\omega \in \Omega$, $K_{\mathbb{U}}(\omega; \Omega) = \mathbb{P}(\Omega) = 1$, hence $\Omega \in \mathcal{C}$. Moreover, for any $\mathbb{S} \in \mathcal{P}(\mathbb{T})$ and any $\omega \in \Omega$, $K_{\mathbb{S}}(\omega; \Omega) = K_{\mathbb{S} \setminus \mathbb{U}}(\omega; \Omega) = 1$, hence $\Omega \in \mathcal{D}$.
- Suppose that $A, B \in \mathcal{C}$ with $A \supseteq B$. Then using the additivity property of probability measures, for any $\omega \in \Omega$,

$$K_{\mathbb{U}}(\omega; A \setminus B) = K_{\mathbb{U}}(\omega; A) - K_{\mathbb{U}}(\omega; B) = \mathbb{P}(A) - \mathbb{P}(B) = \mathbb{P}(A \setminus B),$$

so $A \setminus B \in \mathcal{C}$. If $A, B \in \mathcal{D}$ with $A \supseteq B$, then for any $\mathbb{S} \in \mathcal{P}(\mathbb{T})$ and any $\omega \in \Omega$,

$$\begin{aligned}K_{\mathbb{S}}(\omega; A \setminus B) &= K_{\mathbb{S}}(\omega; A) - K_{\mathbb{S}}(\omega; B) \\ &= K_{\mathbb{S} \setminus \mathbb{U}}(\omega; A) - K_{\mathbb{S} \setminus \mathbb{U}}(\omega; B) \\ &= K_{\mathbb{S} \setminus \mathbb{U}}(\omega; A \setminus B),\end{aligned}$$

so $A \setminus B \in \mathcal{D}$.

- Suppose that $A_n \in \mathcal{C}$ for all $n \in \mathbb{N}$ and $A_1 \subseteq A_2 \subseteq \dots$. Then for any $\omega \in \Omega$, Lemma 2.3.1(iii) tells us that

$$K_{\mathbb{U}}(\omega; \cup_{n \in \mathbb{N}} A_n) = \lim_{n \rightarrow \infty} K_{\mathbb{U}}(\omega; A_n) = \lim_{n \rightarrow \infty} \mathbb{P}(A_n) = \mathbb{P}(\cup_{n \in \mathbb{N}} A_n),$$

so $\cup_{n \in \mathbb{N}} A_n \in \mathcal{C}$. Similarly, if $A_n \in \mathcal{D}$ for all $n \in \mathbb{N}$ with $A_1 \subseteq A_2 \subseteq \dots$, then for any $\mathbb{S} \in \mathcal{P}(\mathbb{T})$ and any $\omega \in \Omega$,

$$K_{\mathbb{S}}(\omega; \cup_{n \in \mathbb{N}} A_n) = \lim_{n \rightarrow \infty} K_{\mathbb{S}}(\omega; A_n) = \lim_{n \rightarrow \infty} K_{\mathbb{S} \setminus \mathbb{U}}(\omega; A_n) = K_{\mathbb{S} \setminus \mathbb{U}}(\omega; \cup_{n \in \mathbb{N}} A_n),$$

so $\cup_{n \in \mathbb{N}} A_n \in \mathcal{D}$.

□

3.3.1 Conditional causal effect

Just as we extend the notion of independence to conditional independence in probability theory, we can also extend the notion of causal effect by conditioning on σ -algebras $\mathcal{G} \subseteq \mathcal{H}$ representing bodies of information. The existence of a conditional causal effect then captures whether or not an intervention via a causal kernel $K_{\mathbb{U}}$ changes the conditional measure given the body of information $\mathcal{G}$ (or equivalently, for the subpopulation that it characterises). For example, we might be interested in the causal effect of a medicine on the blood pressure of patients in a certain age group, rather than averaging out the causal effect across all the age groups.

The extension to the conditional case requires some care, that is not present for the extension of independence to conditional independence. Recall from Section 2.7 that, under a probability measure $\mathbb{P}$, events A and B are said to be conditionally independent given $\mathcal{G}$ if

$$\mathbb{P}_{\mathcal{G}}(\omega, A \cap B) = \mathbb{P}_{\mathcal{G}}(\omega, A)\mathbb{P}_{\mathcal{G}}(\omega, B)$$

for $\mathbb{P}$ -almost all $\omega \in \Omega$. It is this quantifier “for $\mathbb{P}$ -almost all” that requires care in the case of causal effects, since, under interventions, there are multiple probability measures under consideration, which are not necessarily mutually absolutely continuous, rather than a single $\mathbb{P}$.

Unlike the unconditional case, in which causal effect was a binary notion (it either exists or it does not), conditional causal effect is a trichotomy: it can exist, it can be determined not to exist, or it may be undefined whether or not it exists. This is due to the fact that the observational and interventional measures may not be mutually absolutely continuous on the conditioning σ -algebra $\mathcal{G}$. Informally, for an event $A \in \mathcal{H}$,

- if there exists a $\mathcal{G}$ -region in which the probability of A under the observational and interventional measures are mutually absolutely continuous and are different, then we say that there exists a conditional causal effect;
- if the two measures are mutually absolutely continuous everywhere and the same on A , then we say that there does not exist a conditional causal effect;
- otherwise, if the two measures are not mutually absolutely continuous, but in the regions where they are, they are the same, then we say that it is undefined whether or not a conditional causal effect exists.

In order to formalise this idea, we need to recall Lebesgue decompositions (Theorem 2.6.2). For readability, we first present the conditional version of the active causal effect, and we will present the conditional version of the no causal effect later. Take any $\mathbb{U} \in \mathcal{P}(\mathbb{T})$, any outcome $\omega \in \Omega$ and a conditioning σ -algebra $\mathcal{G} \subseteq \mathcal{H}$. Note that $\mathcal{G}$ is not required to be in the form of $\mathcal{H}_{\mathbb{S}}$ for some $\mathbb{S} \in \mathcal{P}(\mathbb{T})$ (see Remark 3.2.6). In this section, we write $\mathbb{P}^{\text{do}(\mathbb{U}, \delta_{\omega})}$ for $K_{\mathbb{U}}(\omega; \cdot)$ for notation convenience, since we will condition this measure on $\mathcal{G}$; it is immediately clear from Definition 3.2.1 that these are the same measure. Let us consider the restrictions of $\mathbb{P}$ and $\mathbb{P}^{\text{do}(\mathbb{U}, \delta_{\omega})}$ on $(\Omega, \mathcal{G})$. Then Theorem 2.6.2 tells us that there exists a measurable subset $\Lambda_{\mathcal{G}}^{\mathbb{U}, \omega} \in \mathcal{G}$ on which $\mathbb{P}$ and $\mathbb{P}^{\text{do}(\mathbb{U}, \delta_{\omega})}$ are mutually

absolutely continuous (the Λ^2 in Theorem 2.6.2), which is maximal in the sense that there exists no $\mathcal{G}$ -measurable subset of $\Omega \setminus \Lambda_{\mathcal{G}}^{\mathbb{U},\omega}$ which has positive measure under both $\mathbb{P}$ and $\mathbb{P}^{\text{do}(\mathbb{U},\delta_\omega)}$ and are mutually absolutely continuous. Moreover, $\Lambda_{\mathcal{G}}^{\mathbb{U},\omega}$ is unique up to $\mathbb{P}$ - and $\mathbb{P}^{\text{do}(\mathbb{U},\delta_\omega)}$ -null sets: if $\tilde{\Lambda} \in \mathcal{G}$ is another measurable subset of Ω with the same properties, then

$$\mathbb{P}(\Lambda_{\mathcal{G}}^{\mathbb{U},\omega} \Delta \tilde{\Lambda}) = \mathbb{P}^{\text{do}(\mathbb{U},\delta_\omega)}(\Lambda_{\mathcal{G}}^{\mathbb{U},\omega} \Delta \tilde{\Lambda}) = 0,$$

where we recall that Δ denotes the symmetric difference between two sets.

Let us define, for any event $A \in \mathcal{H}$, the following set:

$$\Delta_{\mathcal{G}}^{\mathbb{U},\omega}(A) := \left\{ \gamma \in \Lambda_{\mathcal{G}}^{\mathbb{U},\omega} : \mathbb{P}_{\mathcal{G}}^{\text{do}(\mathbb{U},\delta_\omega)}(\gamma, A) \neq \mathbb{P}_{\mathcal{G}}(\gamma, A) \right\}. \quad (3.4)$$

Hence, $\Delta_{\mathcal{G}}^{\mathbb{U},\omega}(A)$ is the set of γ in the overlap region $\Lambda_{\mathcal{G}}^{\mathbb{U},\omega}$ between $\mathbb{P}$ and $\mathbb{P}^{\text{do}(\mathbb{U},\delta_\omega)}$ such that the two $\mathcal{G}$ -measurable random variables

$$\gamma \mapsto \mathbb{P}_{\mathcal{G}}^{\text{do}(\mathbb{U},\delta_\omega)}(\gamma, A) \quad \text{and} \quad \gamma \mapsto \mathbb{P}_{\mathcal{G}}(\gamma, A)$$

differ. Note that both $\mathbb{P}_{\mathcal{G}}^{\text{do}(\mathbb{U},\delta_\omega)}(\cdot, A)$ and $\mathbb{P}_{\mathcal{G}}(\cdot, A)$ are defined up to $\mathbb{P}$ -null sets (since we are only considering $\gamma \in \Lambda_{\mathcal{G}}^{\mathbb{U},\omega}$, this is equivalent to $\mathbb{P}^{\text{do}(\mathbb{U},\delta_\omega)}$ -null sets), so $\Delta_{\mathcal{G}}^{\mathbb{U},\omega}(A)$ itself is only defined up to $\mathbb{P}$ - (or $\mathbb{P}^{\text{do}(\mathbb{U},\delta_\omega)}$ -) null sets. Further, $\Delta_{\mathcal{G}}^{\mathbb{U},\omega}(A)$ belongs to $\mathcal{G}$.

The random variables $\mathbb{P}_{\mathcal{G}}^{\text{do}(\mathbb{U},\delta_\omega)}(\cdot, A)$ and $\mathbb{P}_{\mathcal{G}}(\cdot, A)$ on either side of the equation in the definition of $\Delta_{\mathcal{G}}^{\mathbb{U},\omega}(A)$ are precisely the $\mathcal{G}$ -conditionals of the definition of active causal effects (Definition 3.3.1). Using this, we define the active conditional causal effect as follows.

Definition 3.3.8 (Active conditional causal effect). Take $\mathbb{U} \in \mathcal{P}(\mathbb{T})$, a conditioning σ -algebra $\mathcal{G}$ and an event $A \in \mathcal{H}$. We have the following trichotomy.

- We say that the causal kernel $K_{\mathbb{U}}$ has an active *conditional* causal effect on A given $\mathcal{G}$ if there exists some $\omega \in \Omega$ such that

$$\mathbb{P}(\Delta_{\mathcal{G}}^{\mathbb{U},\omega}(A)) > 0, \quad \text{or equivalently,} \quad \mathbb{P}^{\text{do}(\mathbb{U},\delta_\omega)}(\Delta_{\mathcal{G}}^{\mathbb{U},\omega}(A)) > 0.$$

- We say that $K_{\mathbb{U}}$ does not have an active conditional causal effect on A given $\mathcal{G}$ if, for any $\omega \in \Omega$, the measures $\mathbb{P}$ and $\mathbb{P}^{\text{do}(\mathbb{U},\delta_\omega)}$ are mutually absolutely continuous on $\mathcal{G}$, i.e. $\mathbb{P}(\Lambda_{\mathcal{G}}^{\mathbb{U},\omega}) = \mathbb{P}^{\text{do}(\mathbb{U},\delta_\omega)}(\Lambda_{\mathcal{G}}^{\mathbb{U},\omega}) = 1$, and

$$\mathbb{P}(\Delta_{\mathcal{G}}^{\mathbb{U},\omega}(A)) = \mathbb{P}^{\text{do}(\mathbb{U},\delta_\omega)}(\Delta_{\mathcal{G}}^{\mathbb{U},\omega}(A)) = 0.$$

- Otherwise, it is undefined whether or not $K_{\mathbb{U}}$ has an active conditional causal effect on A given $\mathcal{G}$. This is the case when there are outcomes $\omega \in \Omega$ for which $\mathbb{P}^{\text{do}(\mathbb{U},\delta_\omega)}$ and $\mathbb{P}$ are not mutually absolutely continuous on $\mathcal{G}$, but for all $\omega \in \Omega$, we have

$$\mathbb{P}(\Delta_{\mathcal{G}}^{\mathbb{U},\omega}(A)) = \mathbb{P}^{\text{do}(\mathbb{U},\delta_\omega)}(\Delta_{\mathcal{G}}^{\mathbb{U},\omega}(A)) = 0.$$

Now take a sub- σ -algebra $\mathcal{F} \subseteq \mathcal{H}$.

- We say that K_U has an active conditional causal effect on $\mathcal{F}$ given $\mathcal{G}$ if it has an active conditional causal effect on some $A \in \mathcal{F}$ given $\mathcal{G}$.
- We say that K_U does not have an active conditional causal effect on $\mathcal{F}$ given $\mathcal{G}$ if it does not have an active conditional causal effect on every $A \in \mathcal{F}$ given $\mathcal{G}$.
- Otherwise, we say that it is undefined whether or not K_U has an active conditional causal effect on $\mathcal{F}$ given $\mathcal{G}$. This is the case if, for all $A \in \mathcal{F}$, either K_U does not have an active conditional causal effect on A given $\mathcal{G}$ or it is undefined, and there exists at least one $A \in \mathcal{F}$ for which it is undefined whether or not K_U has an active conditional causal effect on A given $\mathcal{G}$.

The complication in Definition 3.3.8 arises from the fact that conditional measures given $\mathcal{G}$ are only well defined almost surely, coupled with the fact that the observational measure $\mathbb{P}$ and the interventional measure $\mathbb{P}^{\text{do}(U, \delta_\omega)}$ are not guaranteed to be mutually absolutely continuous on $\mathcal{G}$. However, in essence, Definition 3.3.8 is just a direct conditional analogue of active causal effect in Definition 3.3.1; indeed, if the conditioning σ -algebra is trivial, i.e. $\mathcal{G} = \{\emptyset, \Omega\}$, then we recover the statement in Definition 3.3.1. The semantic interpretation is therefore analogous to the unconditional case: the absence of an active conditional causal effect on A given $\mathcal{G}$ means that the observational and interventional regimes are comparable on $\mathcal{G}$, and intervening on the U -coordinates does not change the conditional probability of A away from its observational conditional probability given $\mathcal{G}$.

Let us unpack Definition 3.3.8 a little bit more when $\mathcal{G}$ is not trivial. For each $\omega \in \Omega$, the region we are interested in is $\Lambda_{\mathcal{G}}^{U, \omega}$, in which $\mathbb{P}$ and $\mathbb{P}^{\text{do}(U, \delta_\omega)}$ are mutually absolutely continuous (on $\mathcal{G}$). If there is an event G of positive measure in this overlap region on which the two $\mathcal{G}$ -measurable random variables $\mathbb{P}_{\mathcal{G}}^{\text{do}(U, \delta_\omega)}(\cdot, A)$ and $\mathbb{P}_{\mathcal{G}}(\cdot, A)$ differ, then we conclude that K_U has an active conditional causal effect on A given $\mathcal{G}$. To determine that K_U does not have an active conditional causal effect on A given $\mathcal{G}$, the two measures $\mathbb{P}$ and $\mathbb{P}^{\text{do}(U, \delta_\omega)}$ are required to be mutually absolutely continuous on $\mathcal{G}$, before checking whether $\mathbb{P}_{\mathcal{G}}^{\text{do}(U, \delta_\omega)}(\cdot, A)$ and $\mathbb{P}_{\mathcal{G}}(\cdot, A)$ are the same almost surely. If this were not the case, then in regions with zero measure under one measure and positive measure under the other, the existence of an active conditional causal effect cannot be determined. Note that mutual absolute continuity is only required to hold on $\mathcal{G}$, not on the full σ -algebra $\mathcal{H}$.

In checking for active conditional causal effects, conditioning on $\mathcal{G}$ is done both before and after the intervention. In other words, we condition both $\mathbb{P}$ and $\mathbb{P}^{\text{do}(U, \delta_\omega)}$ on the σ -algebra $\mathcal{G}$. This does not mean that we are conditioning on the same *individuals*—depending on the intervention, it may be that different individuals belong to the subpopulation $\mathcal{G}$ before and after intervention. For discussion of individual-level heterogeneity, we need general causal spaces that take into account heterogeneity, as introduced in Chapter 4.

The intuition behind the conditional causal effect being undefined can be explained with the following simple example. Imagine a study of a drug's effect on survival, where we wish to compare survival probabilities conditioned on $\mathcal{G}$ (e.g. a risk profile). If, under the observational regime, all mass lies in one risk region, while under the interventional regime all mass lies in a different region with no overlap with the observational risk region, then the conditional distributions given $\mathcal{G}$ of $\mathbb{P}$ and $\mathbb{P}^{\text{do}(\mathbb{U}, \delta_\omega)}$ are not simultaneously observable on any common $\mathcal{G}$ -support. In this case the conditional comparison is not well-posed: there is no region of $\mathcal{G}$ on which both measures assign positive mass, so it is meaningless to ask whether the two conditional survival probabilities agree.

Let us see what the axioms imply for active conditional causal effects.

- (i) For any $\omega \in \Omega$, we have $\mathbb{P}^{\text{do}(\emptyset, \delta_\omega)} = \mathbb{P}$, so trivially, for any $A \in \mathcal{H}$ and any $\mathcal{G} \subseteq \mathcal{H}$, we have $\mathbb{P}_{\mathcal{G}}^{\text{do}(\emptyset, \delta_\omega)}(\gamma, A) = \mathbb{P}_{\mathcal{G}}(\gamma, A)$ for $\mathbb{P}$ -almost all $\gamma \in \Omega$, meaning that $K_\emptyset$ does not have an active conditional causal effect on A given $\mathcal{G}$.
- (ii) Now take any $\mathbb{U} \in \mathcal{P}(\mathbb{T})$, any $B \in \mathcal{H}_{\mathbb{U}}$ and any $\omega \in \Omega$. Then by Lemma 3.1.4(ii), we have $\mathbb{P}_{\mathcal{G}}^{\text{do}(\mathbb{U}, \delta_\omega)}(\gamma, B) = \mathbf{1}_B(\omega)$ for $\mathbb{P}^{\text{do}(\mathbb{U}, \delta_\omega)}$ -almost all $\gamma \in \Omega$. Consequently, up to $\mathbb{P}$ - and $\mathbb{P}^{\text{do}(\mathbb{U}, \delta_\omega)}$ -null sets,

$$\Delta_{\mathcal{G}}^{\mathbb{U}, \omega}(B) = \left\{ \gamma \in \Lambda_{\mathcal{G}}^{\mathbb{U}, \omega} : \mathbb{P}_{\mathcal{G}}(\gamma, B) \neq \mathbf{1}_B(\omega) \right\};$$

that is, an active conditional causal effect on B given $\mathcal{G}$ is detected precisely when the observational conditional probability $\mathbb{P}_{\mathcal{G}}(\cdot, B)$ differs, on a non-null part of the overlap region $\Lambda_{\mathcal{G}}^{\mathbb{U}, \omega}$, from the value $\mathbf{1}_B(\omega)$ that the intervention forces.

Suppose now, in addition, that $\mathcal{H}_{\mathbb{U}}$ is non-trivial and that $\mathbb{P}$ and $\mathbb{P}^{\text{do}(\mathbb{U}, \delta_\omega)}$ are mutually absolutely continuous on $\mathcal{G}$ for every $\omega \in \Omega$. Then, exactly as in the unconditional case, $K_{\mathbb{U}}$ has an active conditional causal effect on $\mathcal{H}_{\mathbb{U}}$ given $\mathcal{G}$. Indeed, take $B \in \mathcal{H}_{\mathbb{U}}$ with $B \neq \emptyset$ and $B \neq \Omega$, and pick $\omega \in B$ and $\omega' \in \Omega \setminus B$. If $K_{\mathbb{U}}$ had no active conditional causal effect on B given $\mathcal{G}$, then $\mathbb{P}(\Lambda_{\mathcal{G}}^{\mathbb{U}, \omega}) = \mathbb{P}(\Lambda_{\mathcal{G}}^{\mathbb{U}, \omega'}) = 1$, and the display above would give both $\mathbb{P}_{\mathcal{G}}(\cdot, B) = \mathbf{1}_B(\omega) = 1$ and $\mathbb{P}_{\mathcal{G}}(\cdot, B) = \mathbf{1}_B(\omega') = 0$ $\mathbb{P}$ -almost surely, which is impossible.

Without mutual absolute continuity, however, this conclusion genuinely fails, and it may be undefined whether $K_{\mathbb{U}}$ has an active conditional causal effect on $\mathcal{H}_{\mathbb{U}}$ given $\mathcal{G}$. Take $B \in \mathcal{H}_{\mathbb{U}}$ with $0 < \mathbb{P}(B) < 1$, and condition on $\mathcal{G} = \sigma(B) = \{\emptyset, B, \Omega \setminus B, \Omega\}$, so that $\mathbb{P}_{\mathcal{G}}(\gamma, B) = \mathbf{1}_B(\gamma)$. For $\omega \in B$, we have $\mathbb{P}^{\text{do}(\mathbb{U}, \delta_\omega)}(B) = 1$, so $\mathbb{P}^{\text{do}(\mathbb{U}, \delta_\omega)}$ is carried by B on $\mathcal{G}$, while $\mathbb{P}$ charges both B and $\Omega \setminus B$; hence $\Lambda_{\mathcal{G}}^{\mathbb{U}, \omega} = B$, on which $\mathbb{P}_{\mathcal{G}}(\gamma, B) = 1 = \mathbf{1}_B(\omega)$, and so $\Delta_{\mathcal{G}}^{\mathbb{U}, \omega}(B) = \emptyset$. Symmetrically, for $\omega \in \Omega \setminus B$ we obtain $\Lambda_{\mathcal{G}}^{\mathbb{U}, \omega} = \Omega \setminus B$ and $\Delta_{\mathcal{G}}^{\mathbb{U}, \omega}(B) = \emptyset$. Since $\mathbb{P}(\Lambda_{\mathcal{G}}^{\mathbb{U}, \omega}) < 1$ for every $\omega \in \Omega$, it is undefined whether $K_{\mathbb{U}}$ has an active conditional causal effect on B given $\mathcal{G}$, even though $K_{\mathbb{U}}$ does have an active (unconditional) causal effect on B . The reason is precisely the one anticipated in Definition 3.3.8: here we are conditioning on the very coordinates that are being intervened on, and the observational and interventional regimes are no longer comparable on $\mathcal{G}$.

We can prove the conditional analogue of Theorem 3.3.5(d); in other words, we can characterise the absence of an active conditional causal effect given $\mathcal{G}$ by a conditional independence statement given $\mathcal{G}$ under the interventional probability measure $\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}$, for any assignment measure $\mathbb{Q}$.

Theorem 3.3.9. *Suppose that, for all $\omega \in \Omega$, $\mathbb{P}$ and $\mathbb{P}^{\text{do}(\mathbb{U}, \delta_\omega)}$ are mutually absolutely continuous on $\mathcal{G}$. Then the following three statements are equivalent.*

- (a) $K_{\mathbb{U}}$ does not have an active conditional causal effect on $A \in \mathcal{H}$ given $\mathcal{G}$.
- (b) For any assignment measure $\mathbb{Q}$ on $(\Omega, \mathcal{H}_{\mathbb{U}})$, we have $\mathbb{P}_g^{\text{do}(\mathbb{U}, \mathbb{Q})}(\gamma, A) = \mathbb{P}_g(\gamma, A)$ for $\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}$ -almost all $\gamma \in \Omega$, and for all $B \in \mathcal{H}_{\mathbb{U}}$,

$$\mathbb{P}_g^{\text{do}(\mathbb{U}, \mathbb{Q})}(\gamma, A \cap B) = \mathbb{P}_g^{\text{do}(\mathbb{U}, \mathbb{Q})}(\gamma, A) \mathbb{P}_g^{\text{do}(\mathbb{U}, \mathbb{Q})}(\gamma, B)$$

for $\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}$ -almost all $\gamma \in \Omega$.

- (c) For any assignment measure $\mathbb{Q}$ on $(\Omega, \mathcal{H}_{\mathbb{U}})$ and for $\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}$ -almost all $\gamma \in \Omega$,

$$\mathbb{P}_{g \vee \mathcal{H}_{\mathbb{U}}}^{\text{do}(\mathbb{U}, \mathbb{Q})}(\gamma, A) = \mathbb{P}_g(\gamma, A).$$

Proof. **(a) $\implies$ (b)** First, see that, for any $G \in \mathcal{G}$, and any $B \in \mathcal{H}_{\mathbb{U}}$, using the definition of conditional probabilities and interventional determinism (Theorem 3.1.2(a)),

$$\begin{aligned} \mathbb{E}^{\text{do}(\mathbb{U}, \mathbb{Q})}[\mathbf{1}_G(\cdot) \mathbb{P}_g^{\text{do}(\mathbb{U}, \mathbb{Q})}(\cdot, A \cap B)] &= \mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}(A \cap B \cap G) \\ &= \int \mathbb{Q}(d\omega) K_{\mathbb{U}}(\omega; A \cap B \cap G) \\ &= \int \mathbb{Q}(d\omega) \mathbf{1}_B(\omega) \mathbb{P}^{\text{do}(\mathbb{U}, \delta_\omega)}(A \cap G) \\ &= \int \mathbb{Q}(d\omega) \mathbf{1}_B(\omega) \mathbb{E}^{\text{do}(\mathbb{U}, \delta_\omega)}[\mathbf{1}_G(\cdot) \mathbb{P}_g^{\text{do}(\mathbb{U}, \delta_\omega)}(\cdot, A)]. \end{aligned}$$

Here, we apply (a) to replace $\mathbb{P}_g^{\text{do}(\mathbb{U}, \delta_\omega)}(\cdot, A)$ with $\mathbb{P}_g(\cdot, A)$ inside the expectation to get

$$\begin{aligned} \mathbb{E}^{\text{do}(\mathbb{U}, \mathbb{Q})}[\mathbf{1}_G(\cdot) \mathbb{P}_g^{\text{do}(\mathbb{U}, \mathbb{Q})}(\cdot, A \cap B)] &= \int \mathbb{Q}(d\omega) \mathbf{1}_B(\omega) \mathbb{E}^{\text{do}(\mathbb{U}, \delta_\omega)}[\mathbf{1}_G(\cdot) \mathbb{P}_g(\cdot, A)] \\ &= \int \mathbb{Q}(d\omega) \mathbf{1}_B(\omega) \int K_{\mathbb{U}}(\omega; d\gamma) \mathbf{1}_G(\gamma) \mathbb{P}_g(\gamma, A) \\ &= \int \mathbb{Q}(d\omega) \int K_{\mathbb{U}}(\omega; d\gamma) \mathbf{1}_B(\gamma) \mathbf{1}_G(\gamma) \mathbb{P}_g(\gamma, A), \end{aligned}$$

where, on the last line, we used Theorem 3.1.2(c). Then, since $\mathbf{1}_G(\cdot) \mathbb{P}_g(\cdot, A)$ is $\mathcal{G}$ -measurable, by Proposition 2.7.3(i), we have that

$$\mathbb{E}^{\text{do}(\mathbb{U}, \mathbb{Q})}[\mathbf{1}_G(\cdot) \mathbb{P}_g^{\text{do}(\mathbb{U}, \mathbb{Q})}(\cdot, A \cap B)] = \mathbb{E}^{\text{do}(\mathbb{U}, \mathbb{Q})}[\mathbb{P}_g^{\text{do}(\mathbb{U}, \mathbb{Q})}(\cdot, B) \mathbf{1}_G(\cdot) \mathbb{P}_g(\cdot, A)].$$

Since $G \in \mathcal{G}$ was arbitrary, by Theorem 2.5.3(ii), we have

$$\mathbb{P}_g^{\text{do}(\mathbb{U}, \mathbb{Q})}(\gamma, A \cap B) = \mathbb{P}_g(\gamma, A) \mathbb{P}_g^{\text{do}(\mathbb{U}, \mathbb{Q})}(\gamma, B)$$

for $\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}$ -almost all $\gamma \in \Omega$. Now, letting $B = \Omega$ yields the first statement. Then, using this first statement, we can replace $\mathbb{P}_g(\cdot, A)$ with $\mathbb{P}_g^{\text{do}(\mathbb{U}, \mathbb{Q})}(\cdot, A)$ to obtain

$$\mathbb{P}_g^{\text{do}(\mathbb{U}, \mathbb{Q})}(\cdot, A \cap B) = \mathbb{P}_g^{\text{do}(\mathbb{U}, \mathbb{Q})}(\cdot, A) \mathbb{P}_g^{\text{do}(\mathbb{U}, \mathbb{Q})}(\cdot, B)$$

$\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}$ -almost surely. This means that A and $\mathcal{H}_{\mathbb{U}}$ are conditionally independent given $\mathcal{G}$ under the interventional probability measure $\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}$.

(b) $\implies$ (c) Take any $G \in \mathcal{G}$ and $B \in \mathcal{H}_{\mathbb{U}}$. Then by applying (b), we have

$$\begin{aligned} \mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}(A \cap B \cap G) &= \mathbb{E}^{\text{do}(\mathbb{U}, \mathbb{Q})} \left[\mathbf{1}_G(\cdot) \mathbb{P}_g^{\text{do}(\mathbb{U}, \mathbb{Q})}(\cdot, A \cap B) \right] \\ &= \mathbb{E}^{\text{do}(\mathbb{U}, \mathbb{Q})} \left[\mathbf{1}_G(\cdot) \mathbb{P}_g(\cdot, A) \mathbb{P}_g^{\text{do}(\mathbb{U}, \mathbb{Q})}(\cdot, B) \right]. \end{aligned}$$

Here, since $\mathbf{1}_G(\cdot) \mathbb{P}_g(\cdot, A)$ is $\mathcal{G}$ -measurable, we have, by Proposition 2.7.3(i),

$$\begin{aligned} \mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}(A \cap B \cap G) &= \mathbb{E}^{\text{do}(\mathbb{U}, \mathbb{Q})} [\mathbf{1}_G(\cdot) \mathbb{P}_g(\cdot, A) \mathbf{1}_B(\cdot)] \\ &= \mathbb{E}^{\text{do}(\mathbb{U}, \mathbb{Q})} [\mathbf{1}_{B \cap G}(\cdot) \mathbb{P}_g(\cdot, A)]. \end{aligned}$$

Hence, by Proposition 2.7.3(ii), the random variable $\gamma \mapsto \mathbb{P}_g(\gamma, A)$ is a version of the conditional probability $\gamma \mapsto \mathbb{P}_{g \vee \mathcal{H}_{\mathbb{U}}}^{\text{do}(\mathbb{U}, \mathbb{Q})}(\gamma, A)$.

(c) $\implies$ (a) We will apply (c) with $\mathbb{Q} = \delta_\omega$ for each $\omega \in \Omega$. Take any $G \in \mathcal{G}$. Then since $G \in \mathcal{G} \subseteq \mathcal{G} \vee \mathcal{H}_{\mathbb{U}}$, by the definition of the conditional probability,

$$\mathbb{E}^{\text{do}(\mathbb{U}, \delta_\omega)} [\mathbf{1}_G(\cdot) \mathbb{P}_g^{\text{do}(\mathbb{U}, \delta_\omega)}(\cdot, A)] = \mathbb{P}^{\text{do}(\mathbb{U}, \delta_\omega)}(A \cap G) = \mathbb{E}^{\text{do}(\mathbb{U}, \delta_\omega)} \left[\mathbf{1}_G(\cdot) \mathbb{P}_{g \vee \mathcal{H}_{\mathbb{U}}}^{\text{do}(\mathbb{U}, \delta_\omega)}(\cdot, A) \right].$$

Let us apply (c) here to see that

$$\mathbb{E}^{\text{do}(\mathbb{U}, \delta_\omega)} [\mathbf{1}_G(\cdot) \mathbb{P}_g^{\text{do}(\mathbb{U}, \delta_\omega)}(\cdot, A)] = \mathbb{E}^{\text{do}(\mathbb{U}, \delta_\omega)} [\mathbf{1}_G(\cdot) \mathbb{P}_g(\cdot, A)].$$

Theorem 2.5.3(ii) allows us to conclude immediately. $\square$

Semantically, Theorem 3.3.9(b) says that the intervention does not create any additional dependence between A and the $\mathbb{U}$ -coordinates beyond what is already encoded by $\mathcal{G}$, and Theorem 3.3.9(c) says that, after the intervention, conditioning on the $\mathbb{U}$ -coordinates adds no variation beyond the observational conditional probability given $\mathcal{G}$.

Just as in the unconditional case, the absence of an active conditional causal effect of $K_{\mathbb{U}}$ on A given $\mathcal{G}$ is a sufficient but not a necessary condition to ensure $\mathbb{P}_g^{\text{do}(\mathbb{U}, \mathbb{Q})}(\cdot, A) =$

Outcome	$\omega = (\omega_1, \omega_2)$			
	(0, 0)	(0, 1)	(1, 0)	(1, 1)
$\mathbb{P}$	1/4	1/4	1/4	1/4
$K_{\{1\}}(0; \cdot)$	1/2	1/2	0	0
$K_{\{1\}}(1; \cdot)$	0	0	1/2	1/2

Table 3.4: In Example 3.3.10, $K_{\{1\}}$ has no active causal effect on $A = \{(0, 0), (1, 1)\}$, but has an active conditional causal effect on A given $\mathcal{H}_2$.

$\mathbb{P}_{\mathcal{G}}(\cdot, A)$ $\mathbb{P}$ -almost surely. It is possible that there are $\omega \in \Omega$ such that $\mathbb{P}(\Delta_{\mathcal{G}}^{\mathbb{U}, \omega}(A)) > 0$, but the assignment measure $\mathbb{Q}$ is such that the differences between $\mathbb{P}_{\mathcal{G}}^{\text{do}(\mathbb{U}, \mathbb{Q})}$ and $\mathbb{P}_{\mathcal{G}}$ are cancelled out to give $\mathbb{P}_{\mathcal{G}}^{\text{do}(\mathbb{U}, \mathbb{Q})}(\cdot, A) = \mathbb{P}_{\mathcal{G}}(\cdot, A)$ $\mathbb{P}$ -almost surely.

Theorem 3.3.9(b) was a direct conditional extension of Theorem 3.3.5(i)(d)—indeed, it reduces to Theorem 3.3.5(i)(d) when the conditioning σ -algebra is trivial, i.e. $\mathcal{G} = \{\emptyset, \Omega\}$. Further, Theorem 3.3.9(c) gets absorbed into (b) when $\mathcal{G} = \{\emptyset, \Omega\}$, hence is not interesting when the conditioning σ -algebra is trivial.

Recall from probability theory that, in general, conditional independence does not imply independence, and independence does not imply conditional independence. The same holds for (active) causal effects: even if $K_{\mathbb{U}}$ does not have an active causal effect on an event $A \in \mathcal{H}$, it may have an active conditional causal effect on A given some $\mathcal{G} \subseteq \mathcal{H}$; conversely, $K_{\mathbb{U}}$ may have an active causal effect on A without an active conditional causal effect on A given some σ -algebra $\mathcal{G}$. In the following, we demonstrate a case of this.

Example 3.3.10. For the first direction, take $\mathbb{T} = \{1, 2\}$ with $\Omega_1 = \Omega_2 = \{0, 1\}$ and power-set σ -algebras, let $\mathbb{P}$ be uniform on the four outcomes, and let the causal kernels be as in Table 3.4. Take $\mathbb{U} = \{1\}$, the conditioning σ -algebra $\mathcal{G} = \mathcal{H}_2$, and the “agreement” event $A = \{(0, 0), (1, 1)\}$.

We have $\mathbb{P}(A) = \frac{1}{2}$, and $K_{\{1\}}(0; A) = K_{\{1\}}(1; A) = \frac{1}{2}$, so $K_{\{1\}}$ has no active causal effect on A (Definition 3.3.1). Both $\mathbb{P}$ and $K_{\{1\}}(\omega; \cdot)$ assign measure $\frac{1}{2}$ to each of $\{\omega_2 = 0\}$ and $\{\omega_2 = 1\}$, so $\mathbb{P}$ and $\mathbb{P}^{\text{do}(\{1\}, \delta_{\omega})}$ are mutually absolutely continuous on $\mathcal{G}$ and $\Lambda_{\mathcal{G}}^{\{1\}, \omega} = \Omega$ for every $\omega \in \Omega$. Now

$$\mathbb{P}_{\mathcal{G}}(\gamma, A) = \frac{1}{2} \quad \text{for all } \gamma \in \Omega, \quad \text{whereas} \quad \mathbb{P}_{\mathcal{G}}^{\text{do}(\{1\}, \delta_{\omega})}(\gamma, A) = \mathbf{1}_{\{\gamma_2 = \omega_1\}}(\gamma),$$

since, under $K_{\{1\}}(\omega; \cdot)$, the first coordinate is fixed at ω_1 and so knowing γ_2 determines whether $\gamma \in A$. Hence $\Delta_{\mathcal{G}}^{\{1\}, \omega}(A) = \Omega$ and $\mathbb{P}(\Delta_{\mathcal{G}}^{\{1\}, \omega}(A)) = 1 > 0$, so $K_{\{1\}}$ has an active conditional causal effect on A given $\mathcal{G}$ (Definition 3.3.8). Intervening leaves the probability of agreement untouched, but changes completely how much the second coordinate reveals about it.

For the converse direction, return to the medicine and blood pressure example (Example 3.1.6), and take $\mathbb{U} = \{\text{Med}\}$, $\mathcal{G} = \mathcal{H}_{\text{Blood}}$ and $A = \{\omega_{\text{Blood}} = 0\}$. Since $\mathbb{P}(A) = 0.92$ while $K_{\text{Med}}(0; A) = 0.93$, the kernel $K_{\{\text{Med}\}}$ has an active causal effect on A . But $A \in \mathcal{G}$, so conditional determinism (Proposition 2.7.4) gives $\mathbb{P}_{\mathcal{G}}(\gamma, A) = \mathbf{1}_A(\gamma) =$

$\mathbb{P}_g^{\text{do}(\{\text{Med}\}, \delta_\omega)}(\gamma, A)$ almost surely under either measure, and both measures charge A and $\Omega \setminus A$, so that $\Delta_g^{\{\text{Med}\}, \omega}(A)$ is null in the overlap region. Hence $K_{\{\text{Med}\}}$ has no active conditional causal effect on A given $\mathcal{G}$; this is the instance of Lemma 3.3.13 in which the conditioning σ -algebra already determines the event of interest. $\square$

The following result demonstrates a special case under which the absence of an active conditional causal effect can imply the absence of an active (unconditional) causal effect.

Proposition 3.3.11. *Let $\mathbb{U} \in \mathcal{P}(\mathbb{T})$, $\mathcal{G} \subseteq \mathcal{H}$ and $A \in \mathcal{H}$. Suppose that the following two conditions are satisfied.*

- *A and $\mathcal{G}$ are independent under the observational measure $\mathbb{P}$.*
- *$K_{\mathbb{U}}$ does not have an active conditional causal effect on A given $\mathcal{G}$.*

Then $K_{\mathbb{U}}$ does not have an active causal effect on A .

Proof. Since $K_{\mathbb{U}}$ does not have an active conditional causal effect on A given $\mathcal{G}$, for any $\omega \in \Omega$, the measures $\mathbb{P}^{\text{do}(\mathbb{U}, \delta_\omega)}$ and $\mathbb{P}$ are automatically mutually absolutely continuous on $\mathcal{G}$ (Definition 3.3.8), and for $\mathbb{P}$ -almost all $\gamma \in \Omega$,

$$\mathbb{P}_g^{\text{do}(\mathbb{U}, \delta_\omega)}(\gamma, A) = \mathbb{P}_g(\gamma, A).$$

Since $\mathcal{G}$ and A are independent under $\mathbb{P}$, by Proposition 2.7.5, for $\mathbb{P}$ -almost all $\gamma \in \Omega$, we have $\mathbb{P}_g(\gamma, A) = \mathbb{P}(A)$, so

$$\mathbb{P}_g^{\text{do}(\mathbb{U}, \delta_\omega)}(\gamma, A) = \mathbb{P}(A).$$

Now integrate both sides with respect to $\mathbb{P}^{\text{do}(\mathbb{U}, \delta_\omega)}$ to obtain

$$\mathbb{P}^{\text{do}(\mathbb{U}, \delta_\omega)}(A) = \mathbb{P}(A),$$

which is precisely the definition of no active causal effect of $K_{\mathbb{U}}$ on A , given in Definition 3.3.1. $\square$

The first condition of Proposition 3.3.11 is immediately satisfied if the conditioning σ -algebra is trivial, i.e. $\mathcal{G} = \{\emptyset, \Omega\}$. But then the second condition simply says that $K_{\mathbb{U}}$ does not have an active causal effect on A , so the result is a tautology. Hence, this result is only interesting when the conditioning σ -algebra is not trivial.

Finally, let us turn our attention to the conditional version of no causal effect. Similarly as for $\Lambda_g^{\mathbb{U}, \omega}$, Lebesgue decomposition (Theorem 2.6.2) tells us that, for any $\mathbb{S} \in \mathcal{P}(\mathbb{T})$, there exists a maximal, almost-surely unique measurable subset $\Lambda_g^{\mathbb{S}, \mathbb{S} \setminus \mathbb{U}, \omega} \in \mathcal{G}$ on which $\mathbb{P}^{\text{do}(\mathbb{S}, \delta_\omega)} = K_{\mathbb{S}}(\omega; \cdot)$ and $\mathbb{P}^{\text{do}(\mathbb{S} \setminus \mathbb{U}, \delta_\omega)} = K_{\mathbb{S} \setminus \mathbb{U}}(\omega; \cdot)$ are mutually absolutely continuous. Then, as in Eq. (3.4), let us define, for any event $A \in \mathcal{H}$, the set

$$\Delta_g^{\mathbb{S}, \mathbb{S} \setminus \mathbb{U}, \omega}(A) := \left\{ \gamma \in \Lambda_g^{\mathbb{S}, \mathbb{S} \setminus \mathbb{U}, \omega} : \mathbb{P}_g^{\text{do}(\mathbb{S}, \delta_\omega)}(\gamma, A) \neq \mathbb{P}_g^{\text{do}(\mathbb{S} \setminus \mathbb{U}, \delta_\omega)}(\gamma, A) \right\}. \quad (3.5)$$

Hence, $\Delta_{\mathcal{G}}^{\mathbb{S}, \mathbb{S} \setminus \mathbb{U}, \omega}(A)$ is the set of γ in the overlap region $\Lambda_{\mathcal{G}}^{\mathbb{S}, \mathbb{S} \setminus \mathbb{U}, \omega}$ between $\mathbb{P}^{\text{do}(\mathbb{S}, \delta_{\omega})}$ and $\mathbb{P}^{\text{do}(\mathbb{S} \setminus \mathbb{U}, \delta_{\omega})}$ such that the two $\mathcal{G}$ -measurable functions $\mathbb{P}_{\mathcal{G}}^{\text{do}(\mathbb{S}, \delta_{\omega})}(\cdot, A)$ and $\mathbb{P}_{\mathcal{G}}^{\text{do}(\mathbb{S} \setminus \mathbb{U}, \delta_{\omega})}(\cdot, A)$ differ. Just like $\Delta_{\mathcal{G}}^{\mathbb{U}, \omega}(A)$ in Eq. (3.4), $\Delta_{\mathcal{G}}^{\mathbb{S}, \mathbb{S} \setminus \mathbb{U}, \omega}(A)$ is only defined up to $\mathbb{P}^{\text{do}(\mathbb{S}, \delta_{\omega})}$ - (or equivalently, $\mathbb{P}^{\text{do}(\mathbb{S} \setminus \mathbb{U}, \delta_{\omega})}$ -) null sets, and belongs to $\mathcal{G}$.

We are now ready to define no conditional causal effect, as follows.

Definition 3.3.12 (No conditional causal effect). Take $\mathbb{U} \in \mathcal{P}(\mathbb{T})$, a conditioning σ -algebra $\mathcal{G}$ and an event $A \in \mathcal{H}$. We have the following trichotomy.

- We say that the causal kernel $K_{\mathbb{U}}$ has *no* conditional causal effect on A given $\mathcal{G}$ if, for all $\mathbb{S} \in \mathcal{P}(\mathbb{T})$ and all $\omega \in \Omega$, the measures $\mathbb{P}^{\text{do}(\mathbb{S}, \delta_{\omega})}$ and $\mathbb{P}^{\text{do}(\mathbb{S} \setminus \mathbb{U}, \delta_{\omega})}$ are mutually absolutely continuous on $\mathcal{G}$, i.e. $\mathbb{P}^{\text{do}(\mathbb{S}, \delta_{\omega})}(\Lambda_{\mathcal{G}}^{\mathbb{S}, \mathbb{S} \setminus \mathbb{U}, \omega}) = \mathbb{P}^{\text{do}(\mathbb{S} \setminus \mathbb{U}, \delta_{\omega})}(\Lambda_{\mathcal{G}}^{\mathbb{S}, \mathbb{S} \setminus \mathbb{U}, \omega}) = 1$, and

$$\mathbb{P}^{\text{do}(\mathbb{S}, \delta_{\omega})}(\Delta_{\mathcal{G}}^{\mathbb{S}, \mathbb{S} \setminus \mathbb{U}, \omega}(A)) = \mathbb{P}^{\text{do}(\mathbb{S} \setminus \mathbb{U}, \delta_{\omega})}(\Delta_{\mathcal{G}}^{\mathbb{S}, \mathbb{S} \setminus \mathbb{U}, \omega}(A)) = 0.$$

- We say that $K_{\mathbb{U}}$ has a conditional causal effect on A given $\mathcal{G}$ if there exists some $\mathbb{S} \in \mathcal{P}(\mathbb{T})$ and some $\omega \in \Omega$ such that

$$\mathbb{P}^{\text{do}(\mathbb{S}, \delta_{\omega})}(\Delta_{\mathcal{G}}^{\mathbb{S}, \mathbb{S} \setminus \mathbb{U}, \omega}(A)) > 0, \quad \text{or equivalently,} \quad \mathbb{P}^{\text{do}(\mathbb{S} \setminus \mathbb{U}, \delta_{\omega})}(\Delta_{\mathcal{G}}^{\mathbb{S}, \mathbb{S} \setminus \mathbb{U}, \omega}(A)) > 0.$$

- Otherwise, it is undefined whether or not $K_{\mathbb{U}}$ has a conditional causal effect on A given $\mathcal{G}$. This is the case when there exist $\mathbb{S} \in \mathcal{P}(\mathbb{T})$ and $\omega \in \Omega$ for which $\mathbb{P}^{\text{do}(\mathbb{S}, \delta_{\omega})}$ and $\mathbb{P}^{\text{do}(\mathbb{S} \setminus \mathbb{U}, \delta_{\omega})}$ are not mutually absolutely continuous on $\mathcal{G}$, but for all $\mathbb{S} \in \mathcal{P}(\mathbb{T})$ and all $\omega \in \Omega$, we have

$$\mathbb{P}^{\text{do}(\mathbb{S}, \delta_{\omega})}(\Delta_{\mathcal{G}}^{\mathbb{S}, \mathbb{S} \setminus \mathbb{U}, \omega}(A)) = \mathbb{P}^{\text{do}(\mathbb{S} \setminus \mathbb{U}, \delta_{\omega})}(\Delta_{\mathcal{G}}^{\mathbb{S}, \mathbb{S} \setminus \mathbb{U}, \omega}(A)) = 0.$$

Now take a sub- σ -algebra $\mathcal{F} \subseteq \mathcal{H}$.

- We say that $K_{\mathbb{U}}$ has no conditional causal effect on $\mathcal{F}$ given $\mathcal{G}$ if it has no conditional causal effect on all $A \in \mathcal{F}$ given $\mathcal{G}$.
- We say that $K_{\mathbb{U}}$ has a conditional causal effect on $\mathcal{F}$ given $\mathcal{G}$ if it has a conditional causal effect on some $A \in \mathcal{F}$ given $\mathcal{G}$.
- Otherwise, we say that it is undefined whether or not $K_{\mathbb{U}}$ has a conditional causal effect on $\mathcal{F}$ given $\mathcal{G}$. This is the case when, for all $A \in \mathcal{F}$, either $K_{\mathbb{U}}$ has no conditional causal effect on A given $\mathcal{G}$ or it is undefined, and there exists at least one $A \in \mathcal{F}$ for which it is undefined whether or not $K_{\mathbb{U}}$ has a conditional causal effect on A given $\mathcal{G}$.

Again, the complication in Definition 3.3.12 is due to the fact that conditional measures given $\mathcal{G}$ are only well-defined almost surely, in tandem with the fact that the measures

$\mathbb{P}^{\text{do}(\mathbb{S}, \delta_\omega)}$ and $\mathbb{P}^{\text{do}(\mathbb{S} \setminus \mathbb{U}, \delta_\omega)}$ are not guaranteed to be mutually absolutely continuous on $\mathcal{G}$. The trivial setting of $\mathcal{G} = \{\emptyset, \Omega\}$ allows us to recover the statement of no causal effect in Definition 3.3.1, and for non-trivial $\mathcal{G}$, Definition 3.3.12 is a direct conditional analogue of the corresponding statement in Definition 3.3.1. Unlike the definition of active conditional effect (Definition 3.3.8), we have to check the mutual absolute continuity and the equality of conditionals of $\mathbb{P}^{\text{do}(\mathbb{S}, \delta_\omega)}$ and $\mathbb{P}^{\text{do}(\mathbb{S} \setminus \mathbb{U}, \delta_\omega)}$ for all $\mathbb{S} \in \mathcal{P}(\mathbb{T})$ and all $\omega \in \Omega$.

Just like the unconditional case, if $K_{\mathbb{U}}$ has no conditional causal effect on A given $\mathcal{G}$, then no $K_{\mathbb{S}}$ with $\mathbb{S} \subseteq \mathbb{U}$ can have an active conditional causal effect on A given $\mathcal{G}$. We first have that, for all $\omega \in \Omega$, the measures $\mathbb{P}^{\text{do}(\mathbb{S}, \delta_\omega)}$ and $\mathbb{P}^{\text{do}(\mathbb{S} \setminus \mathbb{U}, \delta_\omega)} = \mathbb{P}^{\text{do}(\emptyset, \delta_\omega)} = \mathbb{P}$ are mutually absolutely continuous on $\mathcal{G}$ by Definition 3.3.12. Then, for all $\omega \in \Omega$, Definition 3.3.12 gives $\mathbb{P}_{\mathcal{G}}^{\text{do}(\mathbb{S}, \delta_\omega)}(\gamma, A) = \mathbb{P}_{\mathcal{G}}(\gamma, A)$ for $\mathbb{P}$ -almost all $\gamma \in \Omega$. These are precisely the two conditions to declare that $K_{\mathbb{U}}$ does not have an active conditional causal effect on A given $\mathcal{G}$ (Definition 3.3.8).

The following is a result that does not have a non-trivial analogue in the unconditional case.

Lemma 3.3.13. *For any $\mathbb{U} \in \mathcal{P}(\mathbb{T})$, the causal kernel $K_{\mathbb{U}}$ either has no conditional causal effect on $\mathcal{G}$ given $\mathcal{G}$, or it is undefined whether or not $K_{\mathbb{U}}$ has a conditional causal effect on $\mathcal{G}$ given $\mathcal{G}$.*

Proof. Take any $\mathbb{S} \in \mathcal{P}(\mathbb{T})$, any $\omega \in \Omega$ and any $G \in \mathcal{G}$. Then by the conditional determinism property (Proposition 2.7.4), we have, for $\mathbb{P}^{\text{do}(\mathbb{S}, \delta_\omega)}$ -almost all $\gamma \in \Omega$,

$$\mathbb{P}_{\mathcal{G}}^{\text{do}(\mathbb{S}, \delta_\omega)}(\gamma, G) = \mathbf{1}_G(\gamma),$$

and for $\mathbb{P}^{\text{do}(\mathbb{S} \setminus \mathbb{U}, \delta_\omega)}$ -almost all $\gamma \in \Omega$,

$$\mathbb{P}_{\mathcal{G}}^{\text{do}(\mathbb{S} \setminus \mathbb{U}, \delta_\omega)}(\gamma, G) = \mathbf{1}_G(\gamma).$$

So for $\mathbb{P}^{\text{do}(\mathbb{S}, \delta_\omega)}$ - (or equivalently, $\mathbb{P}^{\text{do}(\mathbb{S} \setminus \mathbb{U}, \delta_\omega)}$ -)almost all $\gamma \in \Lambda_{\mathcal{G}}^{\mathbb{S}, \mathbb{S} \setminus \mathbb{U}, \omega}$, we have

$$\mathbb{P}_{\mathcal{G}}^{\text{do}(\mathbb{S}, \delta_\omega)}(\gamma, G) = \mathbb{P}_{\mathcal{G}}^{\text{do}(\mathbb{S} \setminus \mathbb{U}, \delta_\omega)}(\gamma, G).$$

Hence, by Definition 3.3.12, if $\mathbb{P}^{\text{do}(\mathbb{S}, \delta_\omega)}$ and $\mathbb{P}^{\text{do}(\mathbb{S} \setminus \mathbb{U}, \delta_\omega)}$ are mutually absolutely continuous on $\mathcal{G}$ for all $\mathbb{S} \in \mathcal{P}(\mathbb{T})$ and $\omega \in \Omega$, then $K_{\mathbb{U}}$ has no conditional causal effect on $\mathcal{G}$ given $\mathcal{G}$, and if there are some $\mathbb{S} \in \mathcal{P}(\mathbb{T})$ and $\omega \in \Omega$ such that they are not mutually absolutely continuous on $\mathcal{G}$, then it is undefined. $\square$

Lemma 3.3.13 should be understood in tandem with the conditional determinism property (Proposition 2.7.4), which says that, once we condition on $\mathcal{G}$, then events in $\mathcal{G}$ are fully determined by the conditioning operation: for all $G \in \mathcal{G}$, we have $\mathbb{P}_{\mathcal{G}}(\gamma, G) = \mathbf{1}_G(\gamma)$ for $\mathbb{P}$ -almost all $\gamma \in \Omega$. Hence, regardless of what intervention takes place on $K_{\mathbb{U}}$, the events $G \in \mathcal{G}$ only listen to the conditioning operation, and so conditional causal effects cannot exist on G . Note that, if $\mathcal{G} = \{\emptyset, \Omega\}$ is trivial, then Lemma 3.3.13 only says that $K_{\mathbb{U}}$ has no (conditional) causal effect on $\emptyset$ and Ω , which is always true—these events always have measures 0 and 1 under any probability measure.

Finally, we define the dormant conditional causal effect.

Definition 3.3.14 (Dormant conditional causal effect). Take $\mathbb{U} \in \mathcal{P}(\mathbb{T})$, a conditioning σ -algebra $\mathcal{G}$ and an event $A \in \mathcal{H}$. We have the following trichotomy.

- We say that $K_{\mathbb{U}}$ has a dormant conditional causal effect on A given $\mathcal{G}$ if there exists some $\mathbb{S} \in \mathcal{P}(\mathbb{T})$ and some $\omega \in \Omega$ such that $\mathbb{P}^{\text{do}(\mathbb{S}, \delta_{\omega})}(\Delta_{\mathcal{G}}^{\mathbb{S}, \mathbb{S} \setminus \mathbb{U}, \omega}(A)) > 0$ (meaning $K_{\mathbb{U}}$ has a conditional causal effect on A given $\mathcal{G}$ —see Definition 3.3.12), but for all $\omega \in \Omega$ such that there exists $\mathbb{S} \in \mathcal{P}(\mathbb{T})$ with $\mathbb{P}^{\text{do}(\mathbb{S}, \delta_{\omega})}(\Delta_{\mathcal{G}}^{\mathbb{S}, \mathbb{S} \setminus \mathbb{U}, \omega}(A)) > 0$, the measures $\mathbb{P}$ and $\mathbb{P}^{\text{do}(\mathbb{U}, \delta_{\omega})}$ are mutually absolutely continuous and $\mathbb{P}(\Delta_{\mathcal{G}}^{\mathbb{U}, \omega}(A)) = 0$.
- If $K_{\mathbb{U}}$ has no conditional causal effect on A given $\mathcal{G}$, or if it has an active conditional causal effect on A given $\mathcal{G}$, then it does not have a dormant conditional causal effect on A given $\mathcal{G}$.
- Otherwise, it is undefined whether or not $K_{\mathbb{U}}$ has a dormant conditional causal effect on A given $\mathcal{G}$.

Now take a sub- σ -algebra $\mathcal{F} \subseteq \mathcal{H}$.

- We say that $K_{\mathbb{U}}$ has a dormant conditional causal effect on $\mathcal{F}$ given $\mathcal{G}$ if it does not have an active conditional causal effect on any $A \in \mathcal{F}$ given $\mathcal{G}$, but it has a conditional causal effect on some $A \in \mathcal{F}$ given $\mathcal{G}$.
- We say that $K_{\mathbb{U}}$ does not have a dormant conditional causal effect on $\mathcal{F}$ given $\mathcal{G}$ if it has either an active conditional causal effect on $\mathcal{F}$ given $\mathcal{G}$ or no conditional causal effect on $\mathcal{F}$ given $\mathcal{G}$.
- Otherwise, we say that it is undefined whether or not $K_{\mathbb{U}}$ has a dormant conditional causal effect on $\mathcal{F}$ given $\mathcal{G}$. This is the case when, for all $A \in \mathcal{F}$, either $K_{\mathbb{U}}$ has no conditional causal effect on A given $\mathcal{G}$ or it is undefined, and there exists at least one $A \in \mathcal{F}$ for which it is undefined whether or not $K_{\mathbb{U}}$ has a conditional causal effect on A given $\mathcal{G}$.

3.3.2 Relative causal effect and causal independence

The conditional extension of causal effects, as considered in Section 3.3.1, has a precise analogue in classical probability theory: the conditional extension of the notion of independence. In this section, we introduce a second, orthogonal extension of causal effects, namely, *relative* causal effects, that has no analogue in classical probability theory.

It does not require such careful consideration as in conditional causal effects, since conditioning is not involved. Instead, relationships between various causal kernels are of interest, which are defined *for all* $\omega \in \Omega$, as discussed in Remark 3.1.3.

Outcome	$\omega = (\omega_1, \omega_2, \omega_3)$							
	(0, 0, 0)	(0, 0, 1)	(0, 1, 0)	(0, 1, 1)	(1, 0, 0)	(1, 0, 1)	(1, 1, 0)	(1, 1, 1)
$\mathbb{P}$	9/32	3/32	1/32	3/32	3/32	1/32	3/32	9/32
$K_1(0; \cdot)$	9/16	3/16	1/16	3/16	0	0	0	0
$K_1(1; \cdot)$	0	0	0	0	3/16	1/16	3/16	9/16
$K_2(0; \cdot)$	3/8	1/8	0	0	3/8	1/8	0	0
$K_{1,2}((0, 0); \cdot)$	3/4	1/4	0	0	0	0	0	0
$K_{1,2}((1, 0); \cdot)$	0	0	0	0	3/4	1/4	0	0

Table 3.5: In Example 3.3.16, K_1 has an active causal effect on $A = \{\omega_3 = 1\}$, but no causal effect on A relative to K_2 .

Definition 3.3.15 (Relative causal effects). Let $\mathbb{U}, \mathbb{W} \in \mathcal{P}(\mathbb{T})$, $A \in \mathcal{H}$ and $\mathcal{F} \subseteq \mathcal{H}$.

- We say that $K_{\mathbb{U}}$ has an active causal effect on A relative to $K_{\mathbb{W}}$ if there exists some $\omega \in \Omega$ such that $K_{\mathbb{U} \cup \mathbb{W}}(\omega; A) \neq K_{\mathbb{W}}(\omega; A)$.

We say that $K_{\mathbb{U}}$ has an active causal effect on $\mathcal{F}$ relative to $K_{\mathbb{W}}$ if it has an active causal effect on some $A \in \mathcal{F}$ relative to $K_{\mathbb{W}}$.

- We say that $K_{\mathbb{U}}$ has no causal effect on A relative to $K_{\mathbb{W}}$ if, for all $\mathbb{S} \in \mathcal{P}(\mathbb{T})$ and all $\omega \in \Omega$, we have $K_{\mathbb{S} \cup \mathbb{W}}(\omega; A) = K_{(\mathbb{S} \cup \mathbb{W}) \setminus (\mathbb{U} \setminus \mathbb{W})}(\omega; A)$.

We say that $K_{\mathbb{U}}$ has no causal effect on $\mathcal{F}$ relative to $K_{\mathbb{W}}$ if it has no causal effect on any $A \in \mathcal{F}$ relative to $K_{\mathbb{W}}$.

- Otherwise, we say that $K_{\mathbb{U}}$ has a dormant causal effect on A relative to $K_{\mathbb{W}}$ (resp. $\mathcal{F}$).

The intuition behind the notion of relative causal effect is as follows. If $K_{\mathbb{U}}$ does not have an active causal effect on A relative to $K_{\mathbb{W}}$, it means that, upon intervention on $K_{\mathbb{W}}$, intervening further on the $\mathbb{U}$ coordinates does not change the measure. The simplest case of this is that of a “causal chain”: if $K_{\mathbb{U}}$ has an active causal effect on A only through $K_{\mathbb{W}}$, then intervening directly on $\mathbb{W}$ will mean that intervening further on $\mathbb{U}$ will not have an active effect. Let us formalise this precisely in the following example.

Example 3.3.16. Let $\mathbb{T} = \{1, 2, 3\}$ with $\Omega_1 = \Omega_2 = \Omega_3 = \{0, 1\}$ and power-set σ -algebras, and consider the mechanism in which each coordinate copies the previous one with probability $\frac{3}{4}$: under $\mathbb{P}$, the first coordinate is uniform, $\mathbb{P}(\omega_2 = \omega_1) = \frac{3}{4}$ and $\mathbb{P}(\omega_3 = \omega_2) = \frac{3}{4}$, with ω_3 depending on ω_1 only through ω_2 . Fixing a coordinate leaves the mechanism downstream of it intact and the coordinates upstream of it at their observational law, giving the causal kernels of Table 3.5. Take $\mathbb{U} = \{1\}$, $\mathbb{W} = \{2\}$ and $A = \{\omega_3 = 1\}$. The marginal observational probability is $\mathbb{P}(A) = \frac{1}{2}$, whereas $K_1(0; A) = \frac{3}{8}$ and $K_1(1; A) = \frac{5}{8}$, so K_1 has an active causal effect on A (Definition 3.3.1). On the other hand,

$$K_{1,2}((0, 0); A) = \frac{1}{4} = K_{1,2}((1, 0); A) = K_2(0; A),$$

and symmetrically $K_{1,2}((\omega_1, 1); A) = \frac{3}{4} = K_2(1; A)$ for both values of ω_1 . Hence $K_{1 \cup 2}(\omega; A) = K_2(\omega; A)$ for every $\omega \in \Omega$, and K_1 does not have an active causal effect on A relative to K_2 (Definition 3.3.15). In fact K_1 has *no* causal effect on A relative to K_2 : for $\mathbb{S} \in \mathcal{P}(\mathbb{T})$ containing 3, both $K_{\mathbb{S} \cup \{2\}}(\omega; A)$ and $K_{(\mathbb{S} \cup \{2\}) \setminus \{1\}}(\omega; A)$ equal $\mathbf{1}_A(\omega)$ by interventional determinism, and the remaining cases are those displayed above.

This is the sense in which the first coordinate acts on the third *only through* the second: once the second coordinate has been fixed by intervention, further intervention on the first changes nothing downstream. $\square$

The intuition behind no relative causal effect is similar: for any $\mathbb{S} \in \mathcal{P}(\mathbb{T})$, intervening on $\mathbb{S} \cup \mathbb{W}$ will be the same as intervening on those coordinates without the $\mathbb{U} \setminus \mathbb{W}$ coordinates. If $K_{\mathbb{U}}$ has no causal effect on A relative to $K_{\mathbb{W}}$, then letting $\mathbb{S} = \mathbb{U}$, we can see that $K_{\mathbb{U} \cup \mathbb{W}}(\omega; A) = K_{(\mathbb{U} \cup \mathbb{W}) \setminus (\mathbb{U} \setminus \mathbb{W})}(\omega; A) = K_{\mathbb{W}}(\omega; A)$, so $K_{\mathbb{U}}$ does not have an active causal effect on A relative to $K_{\mathbb{W}}$.

It is easy to see that all of these definitions reduce directly to the usual causal effects as given in Definition 3.3.1 if $\mathbb{W} = \emptyset$. Here are some simple results that can be derived for relative causal effects.

Lemma 3.3.17. *Let $\mathbb{U}, \mathbb{W} \in \mathcal{P}(\mathbb{T})$ and $A \in \mathcal{H}$.*

- (i) *$K_{\mathbb{U}}$ has no causal effect on $A \in \mathcal{H}$ relative to $K_{\mathbb{U}}$.*
- (ii) *If $K_{\mathbb{U}}$ has no causal effect on A , then for any $\mathbb{W} \in \mathcal{P}(\mathbb{T})$, it has no causal effect on A relative to $K_{\mathbb{W}}$.*
- (iii) *$K_{\mathbb{U}}$ has no causal effect on $\mathcal{H}_{\mathbb{W}}$ relative to $K_{\mathbb{W}}$.*
- (iv) *If $K_{\mathbb{U}}$ has no causal effect on A relative to $K_{\mathbb{W}}$, then for any assignment measure $\mathbb{Q}$ on $(\Omega, \mathcal{H}_{\mathbb{W}})$, the interventional kernel $K_{\mathbb{U} \setminus \mathbb{W}}^{\text{do}(\mathbb{W}, \mathbb{Q})}$ has no causal effect on A .*
- (v) *$K_{\mathbb{U}}$ does not have an active causal effect on A relative to $K_{\mathbb{W}}$ if and only if, for any assignment measure $\mathbb{Q}$ on $(\Omega, \mathcal{H}_{\mathbb{U}})$ and any $\omega \in \Omega$, we have $K_{\mathbb{W}}^{\text{do}(\mathbb{U}, \mathbb{Q})}(\omega; A) = K_{\mathbb{W}}(\omega; A)$.*

Proof. (i) Take any $\mathbb{S} \in \mathcal{P}(\mathbb{T})$. Then for any $\omega \in \Omega$ and any $A \in \mathcal{H}$,

$$K_{(\mathbb{S} \cup \mathbb{U}) \setminus (\mathbb{U} \setminus \mathbb{U})}(\omega; A) = K_{\mathbb{S} \cup \mathbb{U}}(\omega; A).$$

- (ii) Take any $\mathbb{S} \in \mathcal{P}(\mathbb{T})$. Since $K_{\mathbb{U}}$ has no causal effect on A , Lemma 3.3.4(i) tells us that $K_{\mathbb{U} \setminus \mathbb{W}}$ also has no causal effect on A , so

$$K_{\mathbb{S} \cup \mathbb{W}}(\omega; A) = K_{(\mathbb{S} \cup \mathbb{W}) \setminus (\mathbb{U} \setminus \mathbb{W})}(\omega; A).$$

- (iii) Take any $\mathbb{S} \in \mathcal{P}(\mathbb{T})$ and any $B \in \mathcal{H}_{\mathbb{W}}$. Then by applying the interventional determinism axiom twice (since $B \in \mathcal{H}_{\mathbb{W}} \subseteq \mathcal{H}_{(\mathbb{S} \cup \mathbb{W}) \setminus (\mathbb{U} \setminus \mathbb{W})} \subseteq \mathcal{H}_{\mathbb{S} \cup \mathbb{W}}$), we have

$$K_{\mathbb{S} \cup \mathbb{W}}(\omega; B) = \mathbf{1}_B(\omega) = K_{(\mathbb{S} \cup \mathbb{W}) \setminus (\mathbb{U} \setminus \mathbb{W})}(\omega; B).$$

- (iv) Take any $\mathbb{S} \in \mathcal{P}(\mathbb{T})$. Then using the fact that $K_{\mathbb{U}}$ has no causal effect on A relative to $\mathbb{W}$,

$$\begin{aligned} K_{\mathbb{S}}^{\text{do}(\mathbb{W}, \mathbb{Q})}(\omega; A) &= \int \mathbb{Q}(d\omega'_{\mathbb{W} \setminus \mathbb{S}}) K_{\mathbb{S} \cup \mathbb{W}}((\omega_{\mathbb{S}}, \omega'_{\mathbb{W} \setminus \mathbb{S}}); A) \\ &= \int \mathbb{Q}(d\omega'_{(\mathbb{W} \setminus \mathbb{S}) \setminus (\mathbb{U} \setminus \mathbb{W})}) K_{(\mathbb{S} \cup \mathbb{W}) \setminus (\mathbb{U} \setminus \mathbb{W})}((\omega_{\mathbb{S} \setminus (\mathbb{U} \setminus \mathbb{W})}, \omega'_{(\mathbb{W} \setminus \mathbb{S}) \setminus (\mathbb{U} \setminus \mathbb{W})}); A). \end{aligned}$$

Here, we have

$$(\mathbb{S} \cup \mathbb{W}) \setminus (\mathbb{U} \setminus \mathbb{W}) = (\mathbb{S} \setminus (\mathbb{U} \setminus \mathbb{W})) \cup \mathbb{W} \quad \text{and} \quad (\mathbb{W} \setminus \mathbb{S}) \setminus (\mathbb{U} \setminus \mathbb{W}) = \mathbb{W} \setminus (\mathbb{S} \setminus (\mathbb{U} \setminus \mathbb{W})).$$

So

$$\begin{aligned} K_{\mathbb{S}}^{\text{do}(\mathbb{W}, \mathbb{Q})}(\omega; A) &= \int \mathbb{Q}(d\omega'_{\mathbb{W} \setminus (\mathbb{S} \setminus (\mathbb{U} \setminus \mathbb{W}))}) K_{(\mathbb{S} \setminus (\mathbb{U} \setminus \mathbb{W})) \cup \mathbb{W}}((\omega_{\mathbb{S} \setminus (\mathbb{U} \setminus \mathbb{W})}, \omega'_{\mathbb{W} \setminus (\mathbb{S} \setminus (\mathbb{U} \setminus \mathbb{W}))}); A) \\ &= K_{\mathbb{S} \setminus (\mathbb{U} \setminus \mathbb{W})}^{\text{do}(\mathbb{W}, \mathbb{Q})}(\omega; A). \end{aligned}$$

Hence, $K_{\mathbb{U} \setminus \mathbb{W}}^{\text{do}(\mathbb{W}, \mathbb{Q})}$ has no causal effect on A .

- (v) ($\implies$) Take any assignment measure $\mathbb{Q}$ on $(\Omega, \mathcal{H}_{\mathbb{U}})$ and any $\omega \in \Omega$. Then see that

$$\begin{aligned} K_{\mathbb{W}}^{\text{do}(\mathbb{U}, \mathbb{Q})}(\omega; A) &= \int \mathbb{Q}(d\omega'_{\mathbb{U} \setminus \mathbb{W}}) K_{\mathbb{U} \cup \mathbb{W}}((\omega_{\mathbb{W}}, \omega'_{\mathbb{U} \setminus \mathbb{W}}); A) \\ &= \int \mathbb{Q}(d\omega'_{\mathbb{U} \setminus \mathbb{W}}) K_{\mathbb{W}}(\omega; A) \\ &= K_{\mathbb{W}}(\omega; A). \end{aligned}$$

($\impliedby$) See that

$$K_{\mathbb{U} \cup \mathbb{W}}(\omega; A) = K_{\mathbb{W}}^{\text{do}(\mathbb{U}, \delta_{\omega})}(\omega; A) = K_{\mathbb{W}}(\omega; A).$$

□

Lemma 3.3.17(ii) might be surprising—for conditioning, conditional independence does not imply independence nor vice versa, and conditional causal effect does not imply causal effect nor vice versa. With relative causal effect, if $K_{\mathbb{U}}$ has a causal effect on A , it may or may not have a causal effect on A relative to $K_{\mathbb{W}}$, but if $K_{\mathbb{U}}$ has no causal effect on A , it necessarily implies that $K_{\mathbb{U}}$ has no causal effect on A relative to $K_{\mathbb{W}}$.

Now, we want to prove a result that is analogous to Theorem 3.3.5(i) and Theorem 3.3.9, where the absence of an active (conditional) causal effect was shown to be equivalent to a (conditional) independence statement under the interventional measure. The independence statement that we require for relative causal effect is the following.

Definition 3.3.18 (Causal independence). Take $\mathbb{U} \in \mathcal{P}(\mathbb{T})$. We say that a finite number of events $A_1, \dots, A_n \in \mathcal{H}$ are *causally independent* under $K_{\mathbb{U}}$ if, for every

non-empty subset $I \subseteq \{1, \dots, n\}$ and all $\omega \in \Omega$, we have

$$K_{\mathbb{U}}(\omega; \cap_{i \in I} A_i) = \prod_{i \in I} K_{\mathbb{U}}(\omega; A_i).$$

We say that a finite number of sub- σ -algebras $\mathcal{F}_1, \dots, \mathcal{F}_n \subseteq \mathcal{H}$ are causally independent under $K_{\mathbb{U}}$ if all tuples of events $A_1 \in \mathcal{F}_1, \dots, A_n \in \mathcal{F}_n$ are causally independent under $K_{\mathbb{U}}$. Finally, an arbitrary set $\{\mathcal{F}_t\}_{t \in \mathbb{T}}$ of sub- σ -algebras of $\mathcal{H}$, where $\mathbb{T}$ can be uncountably infinite, are said to be causally independent under $K_{\mathbb{U}}$ if every finite subset of them are causally independent under $K_{\mathbb{U}}$.

Note that this is a direct causal analogue of independence ($\mathbb{P}(A \cap B) = \mathbb{P}(A)\mathbb{P}(B)$) and conditional independence ($\mathbb{P}_{\mathcal{G}}(\gamma, A \cap B) = \mathbb{P}_{\mathcal{G}}(\gamma, A)\mathbb{P}_{\mathcal{G}}(\gamma, B)$ for $\mathbb{P}$ -almost all $\gamma \in \Omega$). The subtlety is again that conditional independence is an *almost sure* statement, whereas for causal independence, we need the identity to hold *for all* $\omega \in \Omega$, as discussed in Remark 3.1.3. The intuition is simply that the events A and B are independent upon fixing the $\mathbb{U}$ -coordinates to ω .

We first prove some results on causal independence, then move on to linking active relative causal effects to causal independence.

Lemma 3.3.19. *Take any $\mathbb{U} \in \mathcal{P}(\mathbb{T})$.*

- (i) $\mathcal{H}_{\mathbb{U}}$ is causally independent under $K_{\mathbb{U}}$ with any event $A \in \mathcal{H}$.
- (ii) Events $A, B \in \mathcal{H}$ are causally independent under $K_{\mathbb{U}}$ if and only if, for any assignment measure $\mathbb{Q}$ on $(\Omega, \mathcal{H}_{\mathbb{U}})$, they are conditionally independent under the interventional measure $\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}$ given $\mathcal{H}_{\mathbb{U}}$.

Proof. (i) Take any $B \in \mathcal{H}_{\mathbb{U}}$. Then the interventional determinism axiom (Theorem 3.1.2(a)) immediately gives

$$K_{\mathbb{U}}(\omega; A \cap B) = \mathbf{1}_B(\omega) K_{\mathbb{U}}(\omega; A) = K_{\mathbb{U}}(\omega; B) K_{\mathbb{U}}(\omega; A).$$

- (ii) ($\implies$) Take any assignment measure $\mathbb{Q}$ on $(\Omega, \mathcal{H}_{\mathbb{U}})$. For any $C \in \mathcal{H}_{\mathbb{U}}$, the definition of conditional probabilities and the interventional determinism axiom (Theorem 3.1.2(a)) gives

$$\begin{aligned} \mathbb{E}^{\text{do}(\mathbb{U}, \mathbb{Q})}[\mathbf{1}_C(\cdot) \mathbb{P}_{\mathcal{H}_{\mathbb{U}}}^{\text{do}(\mathbb{U}, \mathbb{Q})}(\cdot, A \cap B)] &= \mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}(A \cap B \cap C) \\ &= \int \mathbb{Q}(d\omega) K_{\mathbb{U}}(\omega; A \cap B \cap C) \\ &= \int \mathbb{Q}(d\omega) \mathbf{1}_C(\omega) K_{\mathbb{U}}(\omega; A \cap B). \end{aligned}$$

Now apply the causal independence hypothesis to see that

$$\mathbb{E}^{\text{do}(\mathbb{U}, \mathbb{Q})}[\mathbf{1}_C(\cdot) \mathbb{P}_{\mathcal{H}_{\mathbb{U}}}^{\text{do}(\mathbb{U}, \mathbb{Q})}(\cdot, A \cap B)] = \int \mathbb{Q}(d\omega) \mathbf{1}_C(\omega) K_{\mathbb{U}}(\omega; A) K_{\mathbb{U}}(\omega; B).$$

Here, the map $\omega \mapsto \mathbf{1}_C(\omega) K_{\mathbb{U}}(\omega; A) K_{\mathbb{U}}(\omega; B)$ is clearly $\mathcal{H}_{\mathbb{U}}$ -measurable, so by Lemma 3.2.2, we have

$$\mathbb{E}^{\text{do}(\mathbb{U}, \mathbb{Q})}[\mathbf{1}_C(\cdot) \mathbb{P}_{\mathcal{H}_{\mathbb{U}}}^{\text{do}(\mathbb{U}, \mathbb{Q})}(\cdot, A \cap B)] = \mathbb{E}^{\text{do}(\mathbb{U}, \mathbb{Q})}[\mathbf{1}_C(\cdot) K_{\mathbb{U}}(\cdot; A) K_{\mathbb{U}}(\cdot; B)].$$

But by the Fundamental Lemma of Causality (Lemma 3.2.7), the kernels $\omega \mapsto K_{\mathbb{U}}(\omega; A)$ and $\omega \mapsto K_{\mathbb{U}}(\omega; B)$ are versions of $\mathbb{P}_{\mathcal{H}_{\mathbb{U}}}^{\text{do}(\mathbb{U}, \mathbb{Q})}(\cdot, A)$ and $\mathbb{P}_{\mathcal{H}_{\mathbb{U}}}^{\text{do}(\mathbb{U}, \mathbb{Q})}(\cdot, B)$ respectively. Hence,

$$\mathbb{E}^{\text{do}(\mathbb{U}, \mathbb{Q})}[\mathbf{1}_C(\cdot) \mathbb{P}_{\mathcal{H}_{\mathbb{U}}}^{\text{do}(\mathbb{U}, \mathbb{Q})}(\cdot, A \cap B)] = \mathbb{E}^{\text{do}(\mathbb{U}, \mathbb{Q})}[\mathbf{1}_C(\cdot) \mathbb{P}_{\mathcal{H}_{\mathbb{U}}}^{\text{do}(\mathbb{U}, \mathbb{Q})}(\cdot, A) \mathbb{P}_{\mathcal{H}_{\mathbb{U}}}^{\text{do}(\mathbb{U}, \mathbb{Q})}(\cdot, B)].$$

Finally, we apply Theorem 2.5.3(ii) to see that

$$\mathbb{P}_{\mathcal{H}_{\mathbb{U}}}^{\text{do}(\mathbb{U}, \mathbb{Q})}(\cdot, A \cap B) = \mathbb{P}_{\mathcal{H}_{\mathbb{U}}}^{\text{do}(\mathbb{U}, \mathbb{Q})}(\cdot, A) \mathbb{P}_{\mathcal{H}_{\mathbb{U}}}^{\text{do}(\mathbb{U}, \mathbb{Q})}(\cdot, B)$$

$\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}$ -almost surely.

($\Leftarrow$) Take any $\omega \in \Omega$. Then the conditional independence hypothesis gives us

$$\mathbb{P}_{\mathcal{H}_{\mathbb{U}}}^{\text{do}(\mathbb{U}, \delta_{\omega})}(\cdot, A \cap B) = \mathbb{P}_{\mathcal{H}_{\mathbb{U}}}^{\text{do}(\mathbb{U}, \delta_{\omega})}(\cdot, A) \mathbb{P}_{\mathcal{H}_{\mathbb{U}}}^{\text{do}(\mathbb{U}, \delta_{\omega})}(\cdot, B).$$

Integrate both sides with respect to $\mathbb{P}^{\text{do}(\mathbb{U}, \delta_{\omega})}$ to see that

$$\mathbb{P}^{\text{do}(\mathbb{U}, \delta_{\omega})}(A \cap B) = \mathbb{E}^{\text{do}(\mathbb{U}, \delta_{\omega})}[\mathbb{P}_{\mathcal{H}_{\mathbb{U}}}^{\text{do}(\mathbb{U}, \delta_{\omega})}(\cdot, A) \mathbb{P}_{\mathcal{H}_{\mathbb{U}}}^{\text{do}(\mathbb{U}, \delta_{\omega})}(\cdot, B)].$$

Here, the left-hand side is the definition of $K_{\mathbb{U}}(\omega; A \cap B)$. On the right-hand side, inside the expectation, we rewrite the conditional probabilities by causal kernels, using the Fundamental Lemma of Causality (Lemma 3.2.7), to obtain

$$\begin{aligned} K_{\mathbb{U}}(\omega; A \cap B) &= \mathbb{E}^{\text{do}(\mathbb{U}, \delta_{\omega})}[K_{\mathbb{U}}(\cdot; A) K_{\mathbb{U}}(\cdot; B)] \\ &= \int K_{\mathbb{U}}(\omega; d\omega') K_{\mathbb{U}}(\omega'; A) K_{\mathbb{U}}(\omega'; B) \\ &= K_{\mathbb{U}}(\omega; A) K_{\mathbb{U}}(\omega; B), \end{aligned}$$

where the third equality follows by applying Theorem 3.1.2(c), since $\omega' \mapsto K_{\mathbb{U}}(\omega'; A) K_{\mathbb{U}}(\omega'; B)$ is $\mathcal{H}_{\mathbb{U}}$ -measurable. $\square$

We finally prove the relative analogue of Theorem 3.3.5(i) and Theorem 3.3.9, linking the notions of active relative causal effect and causal independence. Unlike the previous theorems, this one is a one-way implication, rather than a characterisation. A simpler characterisation is given in Lemma 3.3.17(v).

Theorem 3.3.20. *Take $\mathbb{U}, \mathbb{W} \in \mathcal{P}(\mathbb{T})$, and an event $A \in \mathcal{H}$. If $K_{\mathbb{U}}$ does not have an active causal effect on A relative to $K_{\mathbb{W}}$, then for any assignment measure $\mathbb{Q}$ on $(\Omega, \mathcal{H}_{\mathbb{U}})$ the event A and the σ -algebra $\mathcal{H}_{\mathbb{U}}$ are causally independent under $K_{\mathbb{W}}^{\text{do}(\mathbb{U}, \mathbb{Q})}$.*

Proof. Take any assignment measure $\mathbb{Q}$ on $(\Omega, \mathcal{H}_{\mathbb{U}})$, any $\omega \in \Omega$ and any event $B \in \mathcal{H}_{\mathbb{U}}$. Then by interventional determinism (using $B \in \mathcal{H}_{\mathbb{U}} \subseteq \mathcal{H}_{\mathbb{U}\mathbb{W}}$) and the absence of active causal effect of $K_{\mathbb{U}}$ on A relative to $K_{\mathbb{W}}$,

$$\begin{aligned} K_{\mathbb{W}}^{\text{do}(\mathbb{U}, \mathbb{Q})}(\omega; A \cap B) &= \int \mathbb{Q}(d\omega'_{\mathbb{U}\setminus\mathbb{W}}) K_{\mathbb{U}\mathbb{W}}((\omega_{\mathbb{W}}, \omega'_{\mathbb{U}\setminus\mathbb{W}}); A \cap B) \\ &= \int \mathbb{Q}(d\omega'_{\mathbb{U}\setminus\mathbb{W}}) \mathbf{1}_B((\omega_{\mathbb{W}}, \omega'_{\mathbb{U}\setminus\mathbb{W}})) K_{\mathbb{U}\mathbb{W}}((\omega_{\mathbb{W}}, \omega'_{\mathbb{U}\setminus\mathbb{W}}); A) \\ &= \int \mathbb{Q}(d\omega'_{\mathbb{U}\setminus\mathbb{W}}) K_{\mathbb{U}\mathbb{W}}((\omega_{\mathbb{W}}, \omega'_{\mathbb{U}\setminus\mathbb{W}}); B) K_{\mathbb{W}}(\omega; A) \\ &= K_{\mathbb{W}}^{\text{do}(\mathbb{U}, \mathbb{Q})}(\omega; B) K_{\mathbb{W}}(\omega; A). \end{aligned}$$

Here, $K_{\mathbb{W}}(\omega; A) = K_{\mathbb{W}}^{\text{do}(\mathbb{U}, \mathbb{Q})}(\omega; A)$ by Lemma 3.3.17(v). Hence,

$$K_{\mathbb{W}}^{\text{do}(\mathbb{U}, \mathbb{Q})}(\omega; A \cap B) = K_{\mathbb{W}}^{\text{do}(\mathbb{U}, \mathbb{Q})}(\omega; B) K_{\mathbb{W}}^{\text{do}(\mathbb{U}, \mathbb{Q})}(\omega; A),$$

i.e. A and B are causally independent under $K_{\mathbb{W}}^{\text{do}(\mathbb{U}, \mathbb{Q})}$. $\square$

To end this section, we combine Definitions 3.3.8 and 3.3.15, to define *conditional relative* causal effects. We do not present the complete trichotomy in each case of active, no and dormant causal effects, and instead only give the definitions for the existence and definitive absence of an active causal effect, as well as the case of definitive no causal effect.

Take any $\mathbb{U}, \mathbb{W} \in \mathcal{P}(\mathbb{T})$, $\omega \in \Omega$ and $\mathcal{G} \subseteq \mathcal{H}$. Then define $\Lambda_{\mathcal{G}}^{\mathbb{W}, \mathbb{U}\mathbb{W}, \omega} \in \mathcal{G}$ as the maximal, almost-surely unique set on which $\mathbb{P}^{\text{do}(\mathbb{W}, \delta_{\omega})} = K_{\mathbb{W}}(\omega; \cdot)$ and $\mathbb{P}^{\text{do}(\mathbb{U}\mathbb{W}, \delta_{\omega})} = K_{\mathbb{U}\mathbb{W}}(\omega; \cdot)$ are mutually absolutely continuous on $\mathcal{G}$, that the Lebesgue decomposition theorem (Theorem 2.6.2) gives us. Then, as in Eqs. (3.4) and (3.5), let us define, for any event $A \in \mathcal{H}$, the set

$$\Delta_{\mathcal{G}}^{\mathbb{W}, \mathbb{U}\mathbb{W}, \omega}(A) = \left\{ \gamma \in \Lambda_{\mathcal{G}}^{\mathbb{W}, \mathbb{U}\mathbb{W}, \omega} : \mathbb{P}_{\mathcal{G}}^{\text{do}(\mathbb{W}, \delta_{\omega})}(\gamma, A) \neq \mathbb{P}_{\mathcal{G}}^{\text{do}(\mathbb{U}\mathbb{W}, \delta_{\omega})}(\gamma, A) \right\}.$$

Similarly, for any $\mathbb{S} \in \mathcal{P}(\mathbb{T})$, we also have $\Lambda_{\mathcal{G}}^{\mathbb{S}\mathbb{W}, (\mathbb{S}\mathbb{U}\mathbb{W}) \setminus (\mathbb{U}\setminus\mathbb{W}), \omega} \in \mathcal{G}$, the maximal, almost-surely unique set on which $\mathbb{P}^{\text{do}(\mathbb{S}\mathbb{U}\mathbb{W}, \delta_{\omega})}$ and $\mathbb{P}^{\text{do}((\mathbb{S}\mathbb{U}\mathbb{W}) \setminus (\mathbb{U}\setminus\mathbb{W}), \delta_{\omega})}$ are mutually absolutely continuous on $\mathcal{G}$. Then let us define, for any $A \in \mathcal{H}$, the set $\Delta_{\mathcal{G}}^{\mathbb{S}\mathbb{U}\mathbb{W}, (\mathbb{S}\mathbb{U}\mathbb{W}) \setminus (\mathbb{U}\setminus\mathbb{W}), \omega}(A)$ as

$$\left\{ \gamma \in \Lambda_{\mathcal{G}}^{\mathbb{S}\mathbb{U}\mathbb{W}, (\mathbb{S}\mathbb{U}\mathbb{W}) \setminus (\mathbb{U}\setminus\mathbb{W}), \omega} : \mathbb{P}_{\mathcal{G}}^{\text{do}(\mathbb{S}\mathbb{U}\mathbb{W}, \delta_{\omega})}(\gamma, A) \neq \mathbb{P}_{\mathcal{G}}^{\text{do}((\mathbb{S}\mathbb{U}\mathbb{W}) \setminus (\mathbb{U}\setminus\mathbb{W}), \delta_{\omega})}(\gamma, A) \right\}.$$

Definition 3.3.21 (Conditional relative causal effects). Let $\mathbb{U}, \mathbb{W} \in \mathcal{P}(\mathbb{T})$, $A \in \mathcal{H}$ and $\mathcal{G} \subseteq \mathcal{H}$.

- We say that $K_{\mathbb{U}}$ has an active conditional causal effect on A given $\mathcal{G}$ relative

to $K_{\mathbb{W}}$ if there exists some $\omega \in \Omega$ such that

$$\mathbb{P}^{\text{do}(\mathbb{W}, \delta_\omega)} \left(\Delta_{\mathcal{G}}^{\mathbb{W}, \mathbb{U} \cup \mathbb{W}, \omega}(A) \right) > 0, \text{ equivalently, } \mathbb{P}^{\text{do}(\mathbb{U} \cup \mathbb{W}, \delta_\omega)} \left(\Delta_{\mathcal{G}}^{\mathbb{W}, \mathbb{U} \cup \mathbb{W}, \omega}(A) \right) > 0.$$

- We say that $K_{\mathbb{U}}$ does not have an active conditional causal effect on A given $\mathcal{G}$ relative to $K_{\mathbb{W}}$ if, for all $\omega \in \Omega$, the measures $\mathbb{P}^{\text{do}(\mathbb{W}, \delta_\omega)}$ and $\mathbb{P}^{\text{do}(\mathbb{U} \cup \mathbb{W}, \delta_\omega)}$ are mutually absolutely continuous on $\mathcal{G}$, and

$$\mathbb{P}^{\text{do}(\mathbb{W}, \delta_\omega)} \left(\Delta_{\mathcal{G}}^{\mathbb{W}, \mathbb{U} \cup \mathbb{W}, \omega}(A) \right) = \mathbb{P}^{\text{do}(\mathbb{U} \cup \mathbb{W}, \delta_\omega)} \left(\Delta_{\mathcal{G}}^{\mathbb{W}, \mathbb{U} \cup \mathbb{W}, \omega}(A) \right) = 0.$$

- We say that $K_{\mathbb{U}}$ has no conditional causal effect on A given $\mathcal{G}$ relative to $K_{\mathbb{W}}$ if, for all $\mathbb{S} \in \mathcal{P}(\mathbb{T})$ and all $\omega \in \Omega$, the measures $\mathbb{P}^{\text{do}(\mathbb{S} \cup \mathbb{W}, \delta_\omega)}$ and $\mathbb{P}^{\text{do}((\mathbb{S} \cup \mathbb{W}) \setminus (\mathbb{U} \setminus \mathbb{W}), \delta_\omega)}$ are mutually absolutely continuous on $\mathcal{G}$, and

$$\mathbb{P}^{\text{do}(\mathbb{S} \cup \mathbb{W}, \delta_\omega)} \left(\Delta_{\mathcal{G}}^{\mathbb{S} \cup \mathbb{W}, (\mathbb{S} \cup \mathbb{W}) \setminus (\mathbb{U} \setminus \mathbb{W}), \omega}(A) \right) = 0,$$

or equivalently,

$$\mathbb{P}^{\text{do}((\mathbb{S} \cup \mathbb{W}) \setminus (\mathbb{U} \setminus \mathbb{W}), \delta_\omega)} \left(\Delta_{\mathcal{G}}^{\mathbb{S} \cup \mathbb{W}, (\mathbb{S} \cup \mathbb{W}) \setminus (\mathbb{U} \setminus \mathbb{W}), \omega}(A) \right) = 0.$$

It is easy to see that, if $\mathbb{W} = \emptyset$, then active conditional causal effect on A given $\mathcal{G}$ relative to $K_{\mathbb{W}}$ boils down to conditional causal effect on A given $\mathcal{G}$ (Definition 3.3.8), and no conditional causal effect on A given $\mathcal{G}$ relative to $K_{\mathbb{W}}$ reduces to no conditional causal effect on A given $\mathcal{G}$ (Definition 3.3.12). Likewise, if $\mathcal{G} = \{\emptyset, \Omega\}$ is trivial, then active conditional causal effect on A given $\mathcal{G}$ relative to $K_{\mathbb{W}}$ recovers active causal effect on A relative to $K_{\mathbb{W}}$ (Definition 3.3.15), and no conditional causal effect given $\mathcal{G}$ relative to $K_{\mathbb{W}}$ is the same as no causal effect relative to $K_{\mathbb{W}}$ (Definition 3.3.15).

The natural companion is the notion of *conditional causal independence*, defined as follows.

Definition 3.3.22 (Conditional causal independence). Let $\mathbb{U} \in \mathcal{P}(\mathbb{T})$ and $\mathcal{G} \subseteq \mathcal{H}$. We say that a finite number of events $A_1, \dots, A_n \in \mathcal{H}$ are *causally independent* under $K_{\mathbb{U}}$ conditioned on $\mathcal{G}$ if, for every non-empty subset $I \subseteq \{1, \dots, n\}$, all $\omega \in \Omega$ and $\mathbb{P}^{\text{do}(\mathbb{U}, \delta_\omega)}$ -almost all $\gamma \in \Omega$,

$$\mathbb{P}_{\mathcal{G}}^{\text{do}(\mathbb{U}, \delta_\omega)}(\gamma, \cap_{i \in I} A_i) = \prod_{i \in I} \mathbb{P}_{\mathcal{G}}^{\text{do}(\mathbb{U}, \delta_\omega)}(\gamma, A_i).$$

We say that a finite number of sub- σ -algebras $\mathcal{F}_1, \dots, \mathcal{F}_n \subseteq \mathcal{H}$ are causally independent under $K_{\mathbb{U}}$ conditioned on $\mathcal{G}$ if all tuples of events $A_1 \in \mathcal{F}_1, \dots, A_n \in \mathcal{F}_n$ are causally independent under $K_{\mathbb{U}}$ conditioned on $\mathcal{G}$. Finally, an arbitrary set $\{\mathcal{F}_t\}_{t \in \mathbb{T}}$ of sub- σ -algebras of $\mathcal{H}$, where $\mathbb{T}$ can be uncountably infinite, are said to be causally independent under $K_{\mathbb{U}}$ conditioned on $\mathcal{G}$ if every finite subset of them are so.

Just as Proposition 2.7.5 characterised independence and conditional independence in terms of conditional probabilities, we have the following characterisation of causal independence (Definition 3.3.18) and conditional causal independence (Definition 3.3.22).

Proposition 3.3.23. *Let $\mathbb{U} \in \mathcal{P}(\mathbb{T})$, and let $\mathcal{F}_1, \mathcal{F}_2, \mathcal{G} \subseteq \mathcal{H}$ be sub- σ -algebras.*

- (i) *$\mathcal{F}_1$ and $\mathcal{F}_2$ are causally independent under $K_{\mathbb{U}}$ if and only if, for every $\omega \in \Omega$ and every $A \in \mathcal{F}_2$, we have*

$$\mathbb{P}_{\mathcal{F}_1}^{\text{do}(\mathbb{U}, \delta_\omega)}(\gamma, A) = K_{\mathbb{U}}(\omega; A) \quad \text{for } \mathbb{P}^{\text{do}(\mathbb{U}, \delta_\omega)}\text{-almost all } \gamma \in \Omega.$$

- (ii) *$\mathcal{F}_1$ and $\mathcal{F}_2$ are causally independent under $K_{\mathbb{U}}$ conditioned on $\mathcal{G}$ if and only if, for every $\omega \in \Omega$ and every $A \in \mathcal{F}_2$, we have*

$$\mathbb{P}_{\mathcal{G}}^{\text{do}(\mathbb{U}, \delta_\omega)}(\gamma, A) = \mathbb{P}_{\mathcal{F}_1 \vee \mathcal{G}}^{\text{do}(\mathbb{U}, \delta_\omega)}(\gamma, A) \quad \text{for } \mathbb{P}^{\text{do}(\mathbb{U}, \delta_\omega)}\text{-almost all } \gamma \in \Omega.$$

Proof. Applying Proposition 2.7.5 to the probability space $(\Omega, \mathcal{H}, \mathbb{P}^{\text{do}(\mathbb{U}, \delta_\omega)})$, for each fixed $\omega \in \Omega$ gives both claims. $\square$

Note that we may apply Proposition 3.3.23 to a single event $A \in \mathcal{H}$ by taking $\mathcal{F}_2 = \sigma(A)$: if the displayed identity holds for A , then it holds for $\Omega \setminus A$ by subtraction, and trivially for $\emptyset$ and Ω .

Proposition 3.3.23 also isolates why causal independence required no new measure-theoretic machinery, whereas causal effects and sources did. Causal independence is a statement *within* each interventional measure $\mathbb{P}^{\text{do}(\mathbb{U}, \delta_\omega)}$ separately, so the probabilistic theory of Chapter 2 applies to it verbatim. Causal effects and sources, by contrast, compare quantities *across* measures—the observational $\mathbb{P}$ against the interventional $\mathbb{P}^{\text{do}(\mathbb{U}, \delta_\omega)}$, or one interventional measure against another—and it is exactly there that the interplay of “for all ω ” and “for almost all γ ” of Remark 3.1.3 becomes delicate.

3.4 Sources & Identification

The concept of causal effect, as introduced in the previous section, linked the causal kernels to each other, and to the observational measure. Specifically, the definition of $K_{\mathbb{U}}$ having no causal effect on A linked $K_{\mathbb{S}}$ for all $\mathbb{S} \in \mathcal{P}(\mathbb{T})$ with $K_{\mathbb{S} \setminus \mathbb{U}}$, and the absence of active causal effect linked $K_{\mathbb{U}}$ with the observational measure $\mathbb{P}$. This was extended to the conditional, as well as the post-intervention, or relative, set-up. There is another fundamental way in which they can be related, and that is the content of this section.

Definition 3.4.1 (Sources). We say that the causal kernel $K_{\mathbb{U}}$ is a *source* of an event $A \in \mathcal{H}$ if, for all $B \in \mathcal{H}_{\mathbb{U}}$,

$$\mathbb{P}(A \cap B) = \mathbb{E}[\mathbf{1}_B(\cdot) K_{\mathbb{U}}(\cdot; A)],$$

i.e. $K_{\mathbb{U}}(\gamma; A) = \mathbb{P}_{\mathcal{H}_{\mathbb{U}}}(\gamma, A)$ for $\mathbb{P}$ -almost all $\gamma \in \Omega$.

We say that $K_{\mathbb{U}}$ is a source of a sub- σ -algebra $\mathcal{F} \subseteq \mathcal{H}$ if it is a source of all events $A \in \mathcal{F}$.

In words, $K_{\mathbb{U}}$ is a source of A if the causal kernel $K_{\mathbb{U}}(\cdot; A)$ is a version of the conditional probability $\mathbb{P}_{\mathcal{H}_{\mathbb{U}}}(\cdot, A)$. The intuition is that, $\mathbb{P}$ -almost surely, intervening and conditioning on the $\mathbb{U}$ -coordinates has the same effect on A .

Sources play a crucial role in *identification*; in this context, the problem of expressing a causal quantity in terms of the observational measure. For $\mathbb{U} \neq \emptyset$, the causal kernel $K_{\mathbb{U}}$ is not a priori related to the observational measure $\mathbb{P}$ in any way. Hence, without any assumptions, $K_{\mathbb{U}}$, and the interventional measures $\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}$ for any assignment measure $\mathbb{Q}$, are out of reach if one is only given access to the observational measure $\mathbb{P}$, or indeed, in a statistical setting, access to the data from $\mathbb{P}$.

The absence of an active causal effect from the previous section gives a trivial form of identification: if $K_{\mathbb{U}}$ does not have an active causal effect on A (Definition 3.3.1), then

$$\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}(A) = \int \mathbb{Q}(d\omega) K_{\mathbb{U}}(\omega; A) = \int \mathbb{Q}(d\omega) \mathbb{P}(A) = \mathbb{P}(A),$$

so we have identified the interventional measure $\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}$ exactly with the observational measure.

Sources provide a less trivial but still simple and intuitive way that $\mathbb{P}$ and $K_{\mathbb{U}}$ can be related, that leads to a fundamental identification result.

Proposition 3.4.2 (Identification—sources). *Suppose that $K_{\mathbb{U}}$ is a source of A , and that an assignment measure $\mathbb{Q}$ on $(\Omega, \mathcal{H}_{\mathbb{U}})$ is absolutely continuous with respect to $\mathbb{P}|_{\mathcal{H}_{\mathbb{U}}}$. Then*

$$\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}(A) = \int \mathbb{Q}(d\omega) \mathbb{P}_{\mathcal{H}_{\mathbb{U}}}(\omega, A).$$

Proof. By Theorem 2.5.3(i),

$$\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}(A) = \int \mathbb{Q}(d\omega) K_{\mathbb{U}}(\omega; A) = \int \mathbb{Q}(d\omega) \mathbb{P}_{\mathcal{H}_{\mathbb{U}}}(\omega, A),$$

using the fact that $\mathbb{Q}$ is absolutely continuous with respect to $\mathbb{P}|_{\mathcal{H}_{\mathbb{U}}}$. □

Hence, the interventional measure $\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}$ can be expressed using the observational measure $\mathbb{P}$ and the assignment measure $\mathbb{Q}$ (to which, of course, we have access). Later, in Sections 3.4.1 and 3.4.2, we will explore *conditional* and *relative* variants of sources, that will open door to a wider range of identification options.

As we have done each time we introduced a new concept, let us have a look at what the axioms imply for sources.

- (i) $K_{\emptyset}$ is a source of all events in $\mathcal{H}$. Indeed, the trivial intervention axiom tells us that $K_{\emptyset}(\omega; A) = \mathbb{P}(A)$ for all $\omega \in \Omega$ and $A \in \mathcal{H}$, and conditioning on the trivial σ -algebra $\mathcal{H}_{\emptyset} = \{\emptyset, \Omega\}$ also gives $\mathbb{P}$ back: $\mathbb{P}_{\mathcal{H}_{\emptyset}}(\omega, A) = \mathbb{P}(A)$ for all $\omega \in \Omega$.

- (ii) $K_{\mathbb{U}}$ is always a source of $\mathcal{H}_{\mathbb{U}}$. Indeed, for any $B \in \mathcal{H}_{\mathbb{U}}$, conditional determinism (Proposition 2.7.4) gives $\mathbb{P}_{\mathcal{H}_{\mathbb{U}}}(\omega, B) = \mathbf{1}_B(\omega)$ for $\mathbb{P}$ -almost all $\omega \in \Omega$, and interventional determinism (Axiom (ii)) gives $K_{\mathbb{U}}(\omega; B) = \mathbf{1}_B(\omega)$ for all $\omega \in \Omega$. Hence, we have $K_{\mathbb{U}}(\omega; B) = \mathbb{P}_{\mathcal{H}_{\mathbb{U}}}(\omega, B)$ for $\mathbb{P}$ -almost all $\omega \in \Omega$.

Recall the Fundamental Lemma of Causality (Lemma 3.2.7), which said that $K_{\mathbb{U}}(\cdot; A)$ is a version of the conditional probability $\mathbb{P}_{\mathcal{H}_{\mathbb{U}}}^{\text{do}(\mathbb{U}, \mathbb{Q})}(\cdot, A)$ for any assignment measure $\mathbb{Q}$ on $(\Omega, \mathcal{H}_{\mathbb{U}})$. But we also know from Lemma 3.2.4(i) that $K_{\mathbb{U}}(\omega; A) = K_{\mathbb{U}}^{\text{do}(\mathbb{U}, \mathbb{Q})}(\omega; A)$, for all $\omega \in \Omega$. Hence, it is immediate that $K_{\mathbb{U}}^{\text{do}(\mathbb{U}, \mathbb{Q})}$ is a source of any event A in the interventional causal space $(\Omega, \mathcal{H}, \mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}, \mathbb{K}^{\text{do}(\mathbb{U}, \mathbb{Q})})$.

The following result has a one-line proof, but links the three fundamental concepts of independence, active causal effects (almost) and sources in a beautiful way.

Theorem 3.4.3. *Let us take $\mathbb{U} \in \mathcal{P}(\mathbb{T})$ and an event $A \in \mathcal{H}$. Consider the following three statements.*

- (a) $\mathcal{H}_{\mathbb{U}}$ and A are independent under $\mathbb{P}$.
- (b) $K_{\mathbb{U}}(\omega; A) = \mathbb{P}(A)$ for $\mathbb{P}$ -almost all $\omega \in \Omega$.
- (c) $K_{\mathbb{U}}$ is a source of A .

If any two of the above statements are true, then the third is also true.

Note that (b) is almost the statement that $K_{\mathbb{U}}$ does not have an active causal effect on A , but not quite—it is only stated *almost surely*.

Proof. We have the following characterisations of the three statements:

- (a) $\mathbb{P}_{\mathcal{H}_{\mathbb{U}}}(\omega, A) = \mathbb{P}(A)$ for $\mathbb{P}$ -almost all $\omega \in \Omega$.
- (b) $K_{\mathbb{U}}(\omega; A) = \mathbb{P}(A)$ for $\mathbb{P}$ -almost all $\omega \in \Omega$.
- (c) $K_{\mathbb{U}}(\omega; A) = \mathbb{P}_{\mathcal{H}_{\mathbb{U}}}(\omega, A)$ for $\mathbb{P}$ -almost all $\omega \in \Omega$.

Since $\mathbb{P}$ -almost sure equality is transitive, the claim follows immediately. $\square$

For a fixed $\mathbb{U} \in \mathcal{P}(\mathbb{T})$, the events for which $K_{\mathbb{U}}$ is a source form a λ -system, just as those that $K_{\mathbb{U}}$ has no (active) causal effect did (Lemma 3.3.7).

Lemma 3.4.4. *Let $\mathbb{U} \in \mathcal{P}(\mathbb{T})$. The collection*

$$\mathcal{D} = \{A \in \mathcal{H} : K_{\mathbb{U}}(\omega; A) = \mathbb{P}_{\mathcal{H}_{\mathbb{U}}}(\omega, A) \text{ for } \mathbb{P}\text{-almost all } \omega \in \Omega\}$$

is a λ -system.

Proof. • $K_{\mathbb{U}}(\omega; \Omega) = \mathbb{P}_{\mathcal{H}_{\mathbb{U}}}(\omega, \Omega) = 1$, so $\Omega \in \mathcal{D}$.

- If $A, B \in \mathcal{D}$ with $A \supseteq B$, then for $\mathbb{P}$ -almost all $\omega \in \Omega$, we can apply the linearity of conditional expectations (Lemma 2.7.2(ii)) to see that

$$K_{\mathbb{U}}(\omega; A \setminus B) = K_{\mathbb{U}}(\omega; A) - K_{\mathbb{U}}(\omega; B) = \mathbb{P}_{\mathcal{H}_{\mathbb{U}}}(\omega, A) - \mathbb{P}_{\mathcal{H}_{\mathbb{U}}}(\omega, B) = \mathbb{P}_{\mathcal{H}_{\mathbb{U}}}(\omega, A \setminus B),$$

so $A \setminus B \in \mathcal{D}$.

- If $A_n \in \mathcal{D}$ with $A_1 \subseteq A_2 \subseteq \dots$, then by Lemma 2.3.1(iii) and the monotone convergence property of conditional expectations (Lemma 2.7.2(iii)), for $\mathbb{P}$ -almost all $\omega \in \Omega$,

$$K_{\mathbb{U}}(\omega; \cup_{n \in \mathbb{N}} A_n) = \lim_{n \rightarrow \infty} K_{\mathbb{U}}(\omega; A_n) = \lim_{n \rightarrow \infty} \mathbb{P}_{\mathcal{H}_{\mathbb{U}}}(\omega, A_n) = \mathbb{P}_{\mathcal{H}_{\mathbb{U}}}(\omega, \cup_{n \in \mathbb{N}} A_n),$$

so $\cup_{n \in \mathbb{N}} A_n \in \mathcal{D}$.

□

3.4.1 Conditional sources

Just as independence and causal effect were extended to the conditional setting, we can extend sources to conditional sources. In this section, we again write $\mathbb{P}^{\text{do}(\mathbb{U}, \delta_\omega)}$ for $K_{\mathbb{U}}(\omega; \cdot)$ for notational convenience.

Recall the ordinary definition of sources from Definition 3.4.1: $K_{\mathbb{U}}$ is a source of A if $K_{\mathbb{U}}(\cdot; A)$ is a version of the conditional probability $\mathbb{P}_{\mathcal{H}_{\mathbb{U}}}(\cdot, A)$. An equivalent way of defining sources is as follows. We define a *source version* of $\mathbb{P}_{\mathcal{H}_{\mathbb{U}}}(\cdot, A)$ as a version $\mathbb{P}_{K_{\mathbb{U}}}(\cdot; A)$ that satisfies, for all $\omega \in \Omega$ and all $B \in \mathcal{H}_{\mathbb{U}}$,

$$K_{\mathbb{U}}(\omega; A \cap B) = \int K_{\mathbb{U}}(\omega; d\gamma) \mathbf{1}_B(\gamma) \mathbb{P}_{K_{\mathbb{U}}}(\gamma; A). \quad (3.6)$$

Then $K_{\mathbb{U}}$ is a source of A if and only if a source version of $\mathbb{P}_{\mathcal{H}_{\mathbb{U}}}(\cdot, A)$ exists. Indeed, if $K_{\mathbb{U}}$ is a source of A , then taking $\mathbb{P}_{K_{\mathbb{U}}}(\cdot; A) = K_{\mathbb{U}}(\cdot; A)$ clearly satisfies Eq. (3.6), so a source version exists. Conversely, if a source version $\mathbb{P}_{K_{\mathbb{U}}}(\cdot; A)$ exists, then let $B = \Omega$ in Eq. (3.6). The map $\gamma \mapsto \mathbb{P}_{K_{\mathbb{U}}}(\gamma; A)$ is an $\mathcal{H}_{\mathbb{U}}$ -measurable function, so by interventional determinism (Theorem 3.1.2(c)), we have, for all $\omega \in \Omega$,

$$K_{\mathbb{U}}(\omega; A) = \int K_{\mathbb{U}}(\omega; d\gamma) \mathbb{P}_{K_{\mathbb{U}}}(\gamma; A) = \mathbb{P}_{K_{\mathbb{U}}}(\omega; A),$$

which means that $K_{\mathbb{U}}(\cdot; A)$ is a version of $\mathbb{P}_{\mathcal{H}_{\mathbb{U}}}(\cdot, A)$.

We define conditional sources by conditionalising this equivalent characterisation of sources.

Definition 3.4.5 (Conditional sources). Let $\mathbb{U} \in \mathcal{P}(\mathbb{T})$, $A \in \mathcal{H}$ and $\mathcal{G} \subseteq \mathcal{H}$. We define a *source version* of $\mathbb{P}_{\mathcal{H}_{\mathbb{U} \vee \mathcal{G}}}(\cdot, A)$ as a version $\mathbb{P}_{K_{\mathbb{U}|\mathcal{G}}}(\cdot; A)$ that satisfies, for all

$\omega \in \Omega$, all $B \in \mathcal{H}_{\mathbb{U}}$ and all $G \in \mathcal{G}$,

$$K_{\mathbb{U}}(\omega; A \cap B \cap G) = \int K_{\mathbb{U}}(\omega; d\gamma) \mathbf{1}_{B \cap G}(\gamma) \mathbb{P}_{K_{\mathbb{U}}|\mathcal{G}}(\gamma; A).$$

We say that the causal kernel $K_{\mathbb{U}}$ is a *conditional source* of A given $\mathcal{G}$ if there exists a source version $\mathbb{P}_{K_{\mathbb{U}}|\mathcal{G}}(\cdot; A)$ of $\mathbb{P}_{\mathcal{H}_{\mathbb{U}} \vee \mathcal{G}}(\cdot, A)$.

We say that $K_{\mathbb{U}}$ is a conditional source of a sub- σ -algebra $\mathcal{F} \subseteq \mathcal{H}$ given $\mathcal{G}$ if it is a conditional source of all $A \in \mathcal{F}$ given $\mathcal{G}$.

By Lemma 2.1.3 and Proposition 2.7.3(ii), if $K_{\mathbb{U}}$ is a conditional source of A given $\mathcal{G}$, then a source version $\mathbb{P}_{K_{\mathbb{U}}|\mathcal{G}}(\cdot; A)$ is a version of $\mathbb{P}_{\mathcal{H}_{\mathbb{U}} \vee \mathcal{G}}^{\text{do}(\mathbb{U}, \delta_{\omega})}(\cdot, A) = \mathbb{P}_{\mathcal{G}}^{\text{do}(\mathbb{U}, \delta_{\omega})}(\cdot, A)$ for all $\omega \in \Omega$, where the latter ($\mathbb{P}^{\text{do}(\mathbb{U}, \delta_{\omega})}$ -almost sure) equality follows from Lemma 3.1.4(i). The intuition is precisely the conditional analogue of that of ordinary sources (Definition 3.4.1): after we condition on $\mathcal{G}$, conditioning further on $\mathcal{H}_{\mathbb{U}}$ almost surely has the same effect on A as intervening on $\mathcal{H}_{\mathbb{U}}$.

The careful consideration of source versions is required because $\mathbb{P}_{\mathcal{H}_{\mathbb{U}} \vee \mathcal{G}}(\cdot, A)$ is uniquely defined up to $\mathbb{P}$ -null sets. If $K_{\mathbb{U}}(\omega; \cdot)$ has positive measure on $\mathbb{P}$ -null sets, then the integration on the right-hand side in Definition 3.4.5 would not have been well-defined if we put $\mathbb{P}_{\mathcal{H}_{\mathbb{U}} \vee \mathcal{G}}(\gamma, A)$ in place of $\mathbb{P}_{K_{\mathbb{U}}|\mathcal{G}}(\gamma; A)$. On the other hand, it would be wrong to impose absolute continuity of $K_{\mathbb{U}}(\omega; \cdot)$ with respect to $\mathbb{P}$ as a condition, because, by interventional determinism, this would mean that $\mathbb{P}$ would have to place positive measure on every set in $\mathcal{H}_{\mathbb{U}}$ that is not $\emptyset$. If $\mathcal{G} = \{\emptyset, \Omega\}$ is trivial, we immediately recover Definition 3.4.1, since $\mathcal{H}_{\mathbb{U}} \vee \mathcal{G} = \mathcal{H}_{\mathbb{U}}$ and $B \cap G$ ranges over $\mathcal{H}_{\mathbb{U}}$, so that the defining identity becomes Eq. (3.6). It is worth noting that, in the unconditional case, source versions are uniquely defined, since $K_{\mathbb{U}}(\omega; A)$ is defined *for all* $\omega \in \Omega$ (Remark 3.1.3) and $\mathbb{P}_{K_{\mathbb{U}}}(\omega; A)$ is forced to equal $K_{\mathbb{U}}(\omega; A)$, whereas for the conditional case, source versions are only defined up to sets that have zero measure under $\mathbb{P}$ and $\mathbb{P}^{\text{do}(\mathbb{U}, \delta_{\omega})}$ simultaneously for all $\omega \in \Omega$, since $\mathbb{P}_{K_{\mathbb{U}}|\mathcal{G}}(\cdot; A)$ has to be a version of $\mathbb{P}_{\mathcal{H}_{\mathbb{U}} \vee \mathcal{G}}(\cdot, A)$ and $\mathbb{P}_{\mathcal{H}_{\mathbb{U}} \vee \mathcal{G}}^{\text{do}(\mathbb{U}, \delta_{\omega})}(\cdot, A)$ for every $\omega \in \Omega$.

In the unconditional case, Eq. (3.6) having a solution was automatic—it was the causal kernel itself. In the conditional case, there is no canonical candidate for $\mathbb{P}_{K_{\mathbb{U}}|\mathcal{G}}(\cdot; A)$, and the requirement that a single function serve as the $(\mathcal{H}_{\mathbb{U}} \vee \mathcal{G})$ -conditional probability of A under the observational measure $\mathbb{P}$ and under every interventional measure $\mathbb{P}^{\text{do}(\mathbb{U}, \delta_{\omega})}$ simultaneously is a genuine restriction on the causal space. This is where the interplay between “for all ω ” and “for almost all γ ” of Remark 3.1.3 is at its most delicate.

As seen in the following example, given $\mathcal{G}$, a causal kernel $K_{\mathbb{U}}$ may be a source of A without being a conditional source of A given $\mathcal{G}$, and vice versa.

Example 3.4.6. For a causal kernel that is a source without being a conditional source, take $\mathbb{T} = \{1, 2, 3\}$ with binary coordinates, $\mathbb{U} = \{1\}$, $\mathcal{G} = \mathcal{H}_2$ and $A = \{\omega_3 = 1\}$, and let $\mathbb{P}$ and the causal kernels be as in Table 3.6. The observational conditional probabilities are $\mathbb{P}_{\mathcal{H}_1}(\omega, A) = \frac{1}{2}$ for both values of ω_1 , and $K_1(0; A) = \frac{3}{8} + \frac{1}{8} = \frac{1}{2} = K_1(1; A)$, so K_1 is a source of A (Definition 3.4.1). Conditioning further on the second coordinate,

Outcome	$\omega = (\omega_1, \omega_2, \omega_3)$							
	(0, 0, 0)	(0, 0, 1)	(0, 1, 0)	(0, 1, 1)	(1, 0, 0)	(1, 0, 1)	(1, 1, 0)	(1, 1, 1)
$\mathbb{P}$	3/16	1/16	1/16	3/16	2/16	2/16	2/16	2/16
$K_1(0; \cdot)$	1/8	3/8	3/8	1/8	0	0	0	0
$K_1(1; \cdot)$	0	0	0	0	1/4	1/4	1/4	1/4

Table 3.6: In Example 3.4.6, K_1 is a source of $A = \{\omega_3 = 1\}$, but not a conditional source of A given $\mathcal{H}_2$.

however,

$$\mathbb{P}_{\mathcal{H}_1 \vee \mathcal{G}}(\gamma, A) = \frac{1}{4} \quad \text{and} \quad \mathbb{P}_{\mathcal{H}_1 \vee \mathcal{G}}^{\text{do}(1, \delta_\omega)}(\gamma, A) = \frac{3}{4} \quad \text{for } \gamma = (0, 0, \cdot) \text{ and } \omega_1 = 0,$$

and the outcomes $(0, 0, \cdot)$ have positive measure under both $\mathbb{P}$ and $K_1(0; \cdot)$. No single function can therefore be a version of both, so no source version of $\mathbb{P}_{\mathcal{H}_1 \vee \mathcal{G}}(\cdot, A)$ exists, and K_1 is not a conditional source of A given $\mathcal{G}$ (Definition 3.4.5). Intervening and conditioning on the first coordinate agree on average over the second, but disagree within each of its levels.

For the converse, return to Example 3.1.5, and take $\mathbb{U} = \{1\}$, $\mathcal{G} = \mathcal{H}_2$ and $A = \{\omega_2 = 0\}$. Here $\mathbb{P}_{\mathcal{H}_1}(\omega, A)$ equals 0.8 when $\omega_1 = 0$ and 0.2 when $\omega_1 = 1$, whereas $K_1(0; A) = K_1(1; A) = \frac{1}{2}$, so K_1 is not a source of A . But $A \in \mathcal{G} \subseteq \mathcal{H}_{\mathbb{U}} \vee \mathcal{G}$, so conditional determinism (Proposition 2.7.4) gives $\mathbb{P}_{\mathcal{H}_{\mathbb{U}} \vee \mathcal{G}}(\cdot, A) = \mathbf{1}_A = \mathbb{P}_{\mathcal{H}_{\mathbb{U}} \vee \mathcal{G}}^{\text{do}(1, \delta_\omega)}(\cdot, A)$ almost surely under the relevant measures, and $\mathbf{1}_A$ is a source version. Hence K_1 is a conditional source of A given $\mathcal{G}$. $\square$

The following is an easy result.

Lemma 3.4.7. *For any $\mathbb{U} \in \mathcal{P}(\mathbb{T})$ and any $\mathcal{G} \subseteq \mathcal{H}$, $K_{\mathbb{U}}$ is always a conditional source of $\mathcal{G}$ given $\mathcal{G}$.*

Proof. Take any $G' \in \mathcal{G}$. Then by conditional determinism (Proposition 2.7.4), we can take $\mathbf{1}_{G'}(\cdot)$ as a version of $\mathbb{P}_{\mathcal{H}_{\mathbb{U}} \vee \mathcal{G}}(\cdot, G')$, which then satisfies, for all $\omega \in \Omega$, $B \in \mathcal{H}_{\mathbb{U}}$ and $G \in \mathcal{G}$,

$$\int K_{\mathbb{U}}(\omega; d\gamma) \mathbf{1}_{B \cap G}(\gamma) \mathbf{1}_{G'}(\gamma) = K_{\mathbb{U}}(\omega; G' \cap B \cap G).$$

Hence, $\mathbf{1}_{G'}(\cdot)$ is a source version of $\mathbb{P}_{\mathcal{H}_{\mathbb{U}} \vee \mathcal{G}}(\cdot, G')$. Since $G' \in \mathcal{G}$ was arbitrary, $K_{\mathbb{U}}$ is a conditional source of $\mathcal{G}$ given $\mathcal{G}$. $\square$

Conditional sources broaden the horizon of identification. Let us first state a consequence of Definition 3.4.5.

Proposition 3.4.8 (Identification—conditional sources). *Let $\mathbb{U} \in \mathcal{P}(\mathbb{T})$, $\mathcal{G} \subseteq \mathcal{H}$ and $A \in \mathcal{H}$. Suppose that $\mathbb{Q}$ is an assignment measure on $(\Omega, \mathcal{H}_{\mathbb{U}})$ such that the measure $\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}$ is absolutely continuous with respect to $\mathbb{P}$ on $\mathcal{H}_{\mathbb{U}} \vee \mathcal{G}$. Suppose further that $K_{\mathbb{U}}$ is a conditional source of A given $\mathcal{G}$. Then*

$$\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}(A) = \mathbb{E}^{\text{do}(\mathbb{U}, \mathbb{Q})}[\mathbb{P}_{\mathcal{H}_{\mathbb{U}} \vee \mathcal{G}}(\cdot, A)].$$

Proof. See that, by letting $B = G = \Omega$ in the definition of source versions (Definition 3.4.5),

$$\begin{aligned}\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}(A) &= \int \mathbb{Q}(d\omega) K_{\mathbb{U}}(\omega; A) \\ &= \int \mathbb{Q}(d\omega) \int K_{\mathbb{U}}(\omega; d\gamma) \mathbb{P}_{K_{\mathbb{U}}|\mathcal{G}}(\gamma; A) \\ &= \int \mathbb{Q}(d\omega) \int K_{\mathbb{U}}(\omega; d\gamma) \mathbb{P}_{\mathcal{H}_{\mathbb{U}} \vee \mathcal{G}}(\gamma, A) \\ &= \mathbb{E}^{\text{do}(\mathbb{U}, \mathbb{Q})}[\mathbb{P}_{\mathcal{H}_{\mathbb{U}} \vee \mathcal{G}}(\cdot, A)],\end{aligned}$$

where we could replace the source version $\mathbb{P}_{K_{\mathbb{U}}|\mathcal{G}}(\gamma; A)$ by the conditional probability $\mathbb{P}_{\mathcal{H}_{\mathbb{U}} \vee \mathcal{G}}(\gamma, A)$ by the absolute continuity hypothesis. $\square$

In the unconditional case, $K_{\mathbb{U}}$ being a source of A sufficed to identify the interventional probability $\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}(A)$. However, in the conditional case, the conditional $\gamma \mapsto \mathbb{P}_{\mathcal{H}_{\mathbb{U}} \vee \mathcal{G}}(\gamma, A)$ is $(\mathcal{H}_{\mathbb{U}} \vee \mathcal{G})$ -measurable, and on $\mathcal{H}_{\mathbb{U}} \vee \mathcal{G}$, the interventional measure $\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}$ does not coincide with $\mathbb{Q}$. Hence, we cannot identify $\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}(A)$ without further assumptions that allow us to identify $\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}$ on $\mathcal{H}_{\mathbb{U}} \vee \mathcal{G}$. By Lemmas 2.3.2 and 2.1.3, it suffices to identify $\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}$ on intersections $B \cap G$ with $B \in \mathcal{H}_{\mathbb{U}}$ and $G \in \mathcal{G}$, and by interventional determinism (Theorem 3.1.2(a)), we have

$$\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}(B \cap G) = \int \mathbb{Q}(d\omega) K_{\mathbb{U}}(\omega; B \cap G) = \int \mathbb{Q}(d\omega) \mathbf{1}_B(\omega) K_{\mathbb{U}}(\omega; G).$$

Hence, it suffices to identify $K_{\mathbb{U}}(\omega; G)$ for $\mathbb{Q}$ -almost all $\omega \in \Omega$ and all $G \in \mathcal{G}$. But we already have some identification armoury.

- If $K_{\mathbb{U}}$ does not have an active causal effect on $\mathcal{G}$ (Definition 3.3.1), then $K_{\mathbb{U}}(\omega; G) = \mathbb{P}(G)$ for all $\omega \in \Omega$ and $G \in \mathcal{G}$, so we have identified $K_{\mathbb{U}}(\omega; G)$ with $\mathbb{P}(G)$.
- If $K_{\mathbb{U}}$ is an unconditional source of $\mathcal{G}$ (Definition 3.4.1) and $\mathbb{Q}$ is absolutely continuous with respect to $\mathbb{P}|_{\mathcal{H}_{\mathbb{U}}}$, then $K_{\mathbb{U}}(\omega; G)$ can be identified with $\mathbb{P}_{\mathcal{H}_{\mathbb{U}}}(\omega, G)$ for $\mathbb{Q}$ -almost all $\omega \in \Omega$.

In both of the above cases, $\mathbb{P}^{\text{do}(\mathbb{U}, \mathbb{Q})}(A)$ can then be expressed using $\mathbb{P}$ and $\mathbb{Q}$. It is in this way that conditional sources equip us with a much broader armoury for identification. Even if $K_{\mathbb{U}}$ is not an unconditional source, as long as we can find a conditioning σ -algebra $\mathcal{G}$ such that $K_{\mathbb{U}}$ is a conditional source and which responds to interventions on $\mathcal{H}_{\mathbb{U}}$ in a certain way, we can identify the interventional probability from the observational and assignment measures.

Further, by considering the conditional analogues of the two identification criteria (Definitions 3.3.8 and 3.4.5), identification can also be done in a recursive/sequential manner. Let us formalise this procedure. Recall that it suffices to identify $K_{\mathbb{U}}(\omega; G)$ for $\mathbb{Q}$ -almost all $\omega \in \Omega$ and all $G \in \mathcal{G}$. Let $\mathcal{C} \subseteq \mathcal{H}$ be some other sub- σ -algebra.

- If K_U does not have an active conditional effect on $\mathcal{G}$ given $\mathcal{C}$ (Definition 3.3.8), then

$$K_U(\omega; G) = \mathbb{E}^{\text{do}(U, \delta_\omega)}[\mathbb{P}_{\mathcal{C}}^{\text{do}(U, \delta_\omega)}(\cdot, G)] = \int K_U(\omega; d\gamma) \mathbb{P}_{\mathcal{C}}(\gamma, G),$$

so it now suffices to identify $K_U(\omega; \cdot)$ on $\mathcal{C}$ for $\mathbb{Q}$ -almost all $\omega \in \Omega$, since $\mathbb{P}_{\mathcal{C}}(\cdot, G)$ is $\mathcal{C}$ -measurable.

- If K_U is a conditional source of $\mathcal{G}$ given $\mathcal{C}$ (Definition 3.4.5), then, letting $B = C = \Omega$ in Definition 3.4.5,

$$K_U(\omega; G) = \int K_U(\omega; d\gamma) \mathbb{P}_{K_{U|\mathcal{C}}}(\gamma; G),$$

where the integrand, being a version of $\mathbb{P}_{\mathcal{H}_U \vee \mathcal{C}}(\cdot, G)$, is determined by the observational measure up to $\mathbb{P}$ -null sets. Provided that $\mathbb{P}^{\text{do}(U, \mathbb{Q})}$ is absolutely continuous with respect to $\mathbb{P}$ on $\mathcal{H}_U \vee \mathcal{C}$, so that this ambiguity is immaterial once we average over $\mathbb{Q}$, it now suffices to identify $K_U(\omega; \cdot)$ on $\mathcal{H}_U \vee \mathcal{C}$ for $\mathbb{Q}$ -almost all $\omega \in \Omega$, since $\mathbb{P}_{K_{U|\mathcal{C}}}(\cdot; G)$ is $(\mathcal{H}_U \vee \mathcal{C})$ -measurable.

In both cases, the reduction is in fact to $\mathcal{C}$ alone: by interventional determinism (Theorem 3.1.2(a)), we have $K_U(\omega; B \cap C) = \mathbf{1}_B(\omega) K_U(\omega; C)$ for all $B \in \mathcal{H}_U$ and $C \in \mathcal{C}$, and these intersections form a π -system generating $\mathcal{H}_U \vee \mathcal{C}$ (Lemma 2.1.3), so that Lemma 2.3.2 determines $K_U(\omega; \cdot)$ on $\mathcal{H}_U \vee \mathcal{C}$ from its restriction to $\mathcal{C}$. Each of the two criteria therefore trades the task of identifying $K_U(\omega; \cdot)$ on $\mathcal{G}$ for that of identifying it on $\mathcal{C}$, and the procedure can be iterated.

This yields the following sequential identification scheme. Suppose that $\mathcal{G} = \mathcal{G}_1 \vee \dots \vee \mathcal{G}_n$ for sub- σ -algebras $\mathcal{G}_1, \dots, \mathcal{G}_n \subseteq \mathcal{H}$, and write $\mathcal{G}_{1:i} = \mathcal{G}_1 \vee \dots \vee \mathcal{G}_i$, with $\mathcal{G}_{1:0} = \{\emptyset, \Omega\}$. Suppose that $\mathbb{P}^{\text{do}(U, \mathbb{Q})}$ is absolutely continuous with respect to $\mathbb{P}$ on $\mathcal{H}_U \vee \mathcal{G}$, and that, for each $i = 1, \dots, n$, the causal kernel K_U is a conditional source of $\mathcal{G}_i$ given $\mathcal{G}_{1:i-1}$; for $i = 1$ this is an unconditional source of $\mathcal{G}_1$, by the remark following Definition 3.4.5. Then $\mathbb{P}^{\text{do}(U, \mathbb{Q})}(A)$ is identified.

Indeed, we show by induction on i that $K_U(\omega; \cdot)$ is identified on $\mathcal{H}_U \vee \mathcal{G}_{1:i}$ for $\mathbb{Q}$ -almost all $\omega \in \Omega$. For $i = 0$ this is interventional determinism. Suppose it holds for $i - 1$, and take $G \in \mathcal{G}_i$, $B \in \mathcal{H}_U$ and $C \in \mathcal{G}_{1:i-1}$. Applying Definition 3.4.5 with the conditioning σ -algebra $\mathcal{G}_{1:i-1}$,

$$K_U(\omega; G \cap B \cap C) = \int K_U(\omega; d\gamma) \mathbf{1}_{B \cap C}(\gamma) \mathbb{P}_{K_{U|\mathcal{G}_{1:i-1}}}(\gamma; G).$$

The integrand is $(\mathcal{H}_U \vee \mathcal{G}_{1:i-1})$ -measurable and, up to $\mathbb{P}$ -null sets, determined by $\mathbb{P}$; the measure it is integrated against is identified on $\mathcal{H}_U \vee \mathcal{G}_{1:i-1}$ by the induction hypothesis. Hence the right-hand side is identified. Since the sets $G \cap B \cap C$ form a π -system generating $\mathcal{H}_U \vee \mathcal{G}_{1:i}$ and containing Ω , Lemma 2.3.2 identifies $K_U(\omega; \cdot)$ on $\mathcal{H}_U \vee \mathcal{G}_{1:i}$, completing the induction.

Taking $i = n$ identifies $\mathbb{P}^{\text{do}(U, \mathbb{Q})}$ on $\mathcal{H}_U \vee \mathcal{G}$, and Proposition 3.4.8 then identifies $\mathbb{P}^{\text{do}(U, \mathbb{Q})}(A)$.

3.4.2 Relative sources

Finally, we introduce the relative variant of sources.

Definition 3.4.9 (Relative sources). We say that the causal kernel $K_{\mathbb{U}}$ is a source of an event $A \in \mathcal{H}$ *relative* to $K_{\mathbb{W}}$ if, for all $\omega \in \Omega$ and all $B \in \mathcal{H}_{\mathbb{U}}$,

$$K_{\mathbb{W}}(\omega; A \cap B) = \int K_{\mathbb{W}}(\omega; d\gamma) \mathbf{1}_B(\gamma) K_{\mathbb{U}}^{\text{do}(\mathbb{W}, \delta_\omega)}(\gamma; A).$$

We say that $K_{\mathbb{U}}$ is a source of a sub- σ -algebra $\mathcal{F}$ relative to $K_{\mathbb{W}}$ if it is a source of all $A \in \mathcal{F}$ relative to $K_{\mathbb{W}}$.

Mathematically, this means that $K_{\mathbb{U}}^{\text{do}(\mathbb{W}, \delta_\omega)}(\cdot; A)$ is a version of the conditional probability $\mathbb{P}_{\mathcal{H}_{\mathbb{U}}}^{\text{do}(\mathbb{W}, \delta_\omega)}(\cdot, A)$. The intuition is analogous to the previous two definitions: once we intervene on $\mathcal{H}_{\mathbb{W}}$, intervening further on $\mathcal{H}_{\mathbb{U}}$ is the same as conditioning on $\mathcal{H}_{\mathbb{U}}$. If $\mathbb{W} = \emptyset$, then we immediately recover the ordinary sources in Definition 3.4.1.

Recall that there were two possible directions of defining sources: the first is to directly require $K_{\mathbb{U}}(\cdot; A)$ to be a version of $\mathbb{P}_{\mathcal{H}_{\mathbb{U}}}$, and the second is to ask for the existence of a source version. The first way is not amenable to conditionalisation, so for Definition 3.4.5, we conditionalised the second route. However, there is no doubt that the first is the more intuitive, and that is the route that we take in Definition 3.4.9 to define relative sources.

Using the definition of $K_{\mathbb{U}}^{\text{do}(\mathbb{W}, \delta_\omega)}(\cdot; A)$ (Eq. (3.2)), we also have the following equivalent characterisation of relative sources: $K_{\mathbb{U}}$ is a source of A relative to $K_{\mathbb{W}}$ if, for all $\omega \in \Omega$ and all $B \in \mathcal{H}_{\mathbb{U}}$,

$$K_{\mathbb{W}}(\omega; A \cap B) = \int K_{\mathbb{W}}(\omega; d\gamma) \mathbf{1}_B(\gamma) K_{\mathbb{U} \cup \mathbb{W}}((\gamma_{\mathbb{U}}, \omega_{\mathbb{W} \setminus \mathbb{U}}); A).$$

Relative sources allow for further identification results, but of a different kind. So far, identification meant identifying an interventional quantity from the observational measure $\mathbb{P}$ and the assignment measure $\mathbb{Q}$. Here, we will be identifying an interventional measure with another interventional measure, where the intervention takes place on a different set of coordinates.

Proposition 3.4.10 (Identification—relative sources). *Take $\mathbb{U}, \mathbb{W} \in \mathcal{P}(\mathbb{T})$ and $A \in \mathcal{H}$, and let $\mathbb{Q}_{\mathbb{U}}$ and $\mathbb{Q}_{\mathbb{W}}$ be assignment measures on $(\Omega, \mathcal{H}_{\mathbb{U}})$ and $(\Omega, \mathcal{H}_{\mathbb{W} \setminus \mathbb{U}})$ respectively. Suppose that $K_{\mathbb{U}}$ is a source of A relative to $K_{\mathbb{W}}$, and that, for $\mathbb{Q}_{\mathbb{W}}$ -almost all $\eta \in \Omega$, the assignment measure $\mathbb{Q}_{\mathbb{U}}$ is absolutely continuous with respect to $\mathbb{P}^{\text{do}(\mathbb{W}, \delta_\eta)}$ on $\mathcal{H}_{\mathbb{U}}$. Then*

$$\mathbb{P}^{\text{do}(\mathbb{U} \cup \mathbb{W}, \mathbb{Q}_{\mathbb{U}} \otimes \mathbb{Q}_{\mathbb{W}})}(A) = \int \mathbb{Q}_{\mathbb{W}}(d\eta) \int \mathbb{Q}_{\mathbb{U}}(d\gamma) \mathbb{P}_{\mathcal{H}_{\mathbb{U}}}^{\text{do}(\mathbb{W}, \delta_\eta)}(\gamma, A).$$

Proof. See that, by Eq. (3.1) and Fubini's theorem (Theorem 2.5.6),

$$\mathbb{P}^{\text{do}(\mathbb{U} \cup \mathbb{W}, \mathbb{Q}_{\mathbb{U}} \otimes \mathbb{Q}_{\mathbb{W}})}(A) = \int \mathbb{Q}_{\mathbb{U}} \otimes \mathbb{Q}_{\mathbb{W}}(d\omega) K_{\mathbb{U} \cup \mathbb{W}}(\omega; A)$$

$$\begin{aligned}
 &= \int \mathbb{Q}_{\mathbb{W}}(d\eta) \int \mathbb{Q}_{\mathbb{U}}(d\gamma) K_{\mathbb{U}\cup\mathbb{W}}((\gamma_{\mathbb{U}}, \eta_{\mathbb{W}\setminus\mathbb{U}}); A) \\
 &= \int \mathbb{Q}_{\mathbb{W}}(d\eta) \int \mathbb{Q}_{\mathbb{U}}(d\gamma) K_{\mathbb{U}}^{\text{do}(\mathbb{W}, \delta_{\eta})}(\gamma; A),
 \end{aligned}$$

where the last equality is the definition of the interventional causal kernels (Eq. (3.2)), since $\int \delta_{\eta}(d\gamma'_{\mathbb{W}\setminus\mathbb{U}}) K_{\mathbb{U}\cup\mathbb{W}}((\gamma_{\mathbb{U}}, \gamma'_{\mathbb{W}\setminus\mathbb{U}}); A) = K_{\mathbb{U}\cup\mathbb{W}}((\gamma_{\mathbb{U}}, \eta_{\mathbb{W}\setminus\mathbb{U}}); A)$. Now, $K_{\mathbb{U}}$ being a source of A relative to $K_{\mathbb{W}}$ means precisely that $K_{\mathbb{U}}^{\text{do}(\mathbb{W}, \delta_{\eta})}(\cdot; A)$ is a version of the conditional probability $\mathbb{P}_{\mathcal{H}_{\mathbb{U}}}^{\text{do}(\mathbb{W}, \delta_{\eta})}(\cdot, A)$, for every $\eta \in \Omega$ (Definition 3.4.9). Since $\mathbb{Q}_{\mathbb{U}}$ is absolutely continuous with respect to $\mathbb{P}^{\text{do}(\mathbb{W}, \delta_{\eta})}$ on $\mathcal{H}_{\mathbb{U}}$ for $\mathbb{Q}_{\mathbb{W}}$ -almost all $\eta \in \Omega$, the choice of version does not affect the inner integral, by Theorem 2.5.3(i), and the claim follows. $\square$

By Lemma 3.2.4(ii), the left-hand side is the measure obtained by intervening first on $\mathcal{H}_{\mathbb{W}}$ with $\mathbb{Q}_{\mathbb{W}}$ and then on $\mathcal{H}_{\mathbb{U}}$ with $\mathbb{Q}_{\mathbb{U}}$, so Proposition 3.4.10 says that the effect of the second intervention can be read off from the first: having performed the $\mathbb{W}$ -intervention, we may simply *condition* on the $\mathbb{U}$ -coordinates within the resulting interventional measure.

If, in addition, $K_{\mathbb{W}}$ is a source of $\mathcal{H}_{\mathbb{U}} \vee \sigma(A)$ and $\mathbb{Q}_{\mathbb{W}}$ is absolutely continuous with respect to $\mathbb{P}$, then $\mathbb{P}^{\text{do}(\mathbb{W}, \delta_{\eta})}$ and $\mathbb{P}_{\mathcal{H}_{\mathbb{W}}}(\eta, \cdot)$ agree on $\mathcal{H}_{\mathbb{U}} \vee \sigma(A)$ for $\mathbb{Q}_{\mathbb{W}}$ -almost all $\eta \in \Omega$, and the conditional probability $\mathbb{P}_{\mathcal{H}_{\mathbb{U}}}^{\text{do}(\mathbb{W}, \delta_{\eta})}(\cdot, A)$, which depends only on this restriction, is itself identified from $\mathbb{P}$. In that case $\mathbb{P}^{\text{do}(\mathbb{U}\cup\mathbb{W}, \mathbb{Q}_{\mathbb{U}}\otimes\mathbb{Q}_{\mathbb{W}})}(A)$ can be expressed using $\mathbb{P}$, $\mathbb{Q}_{\mathbb{U}}$ and $\mathbb{Q}_{\mathbb{W}}$ alone, and we obtain a two-stage identification formula in which the $\mathbb{W}$ -stage is handled by an ordinary source and the $\mathbb{U}$ -stage by a relative one.

Let us state the relative analogue of Theorem 3.4.3. Recall that the independence of $\mathcal{H}_{\mathbb{U}}$ and A , the almost-sure absence of an active causal effect, and $K_{\mathbb{U}}$ being a source of A were three $\mathbb{P}$ -almost sure identities among the causal kernel $K_{\mathbb{U}}(\cdot; A)$, the conditional probability $\mathbb{P}_{\mathcal{H}_{\mathbb{U}}}(\cdot, A)$ and the constant $\mathbb{P}(A)$. The relative version replaces the observational measure $\mathbb{P}$ throughout by the post-intervention measure $\mathbb{P}^{\text{do}(\mathbb{W}, \delta_{\omega})}$.

Theorem 3.4.11. *Let $\mathbb{U}, \mathbb{W} \in \mathcal{P}(\mathbb{T})$ and $A \in \mathcal{H}$. Consider the following three statements.*

- (i) $K_{\mathbb{U}}$ is a source of A relative to $K_{\mathbb{W}}$.
- (ii) For all $\omega \in \Omega$, $K_{\mathbb{U}}^{\text{do}(\mathbb{W}, \delta_{\omega})}(\gamma; A) = K_{\mathbb{W}}(\omega; A)$ for $\mathbb{P}^{\text{do}(\mathbb{W}, \delta_{\omega})}$ -almost all $\gamma \in \Omega$.
- (iii) $\mathcal{H}_{\mathbb{U}}$ and A are causally independent under $K_{\mathbb{W}}$.

If any two of the three above statements are true, then the third one is also true.

Proof. Let us write down the mathematical statements. The second statement was already given mathematically, and is written verbatim.

- (i) For all $\omega \in \Omega$, $K_{\mathbb{U}}^{\text{do}(\mathbb{W}, \delta_{\omega})}(\gamma; A) = \mathbb{P}_{\mathcal{H}_{\mathbb{U}}}^{\text{do}(\mathbb{W}, \delta_{\omega})}(\gamma, A)$ for $\mathbb{P}^{\text{do}(\mathbb{W}, \delta_{\omega})}$ -almost all $\gamma \in \Omega$.

- (ii) For all $\omega \in \Omega$, $K_{\mathbb{U}}^{\text{do}(\mathbb{W}, \delta_\omega)}(\gamma; A) = K_{\mathbb{W}}(\omega; A)$ for $\mathbb{P}^{\text{do}(\mathbb{W}, \delta_\omega)}$ -almost all $\gamma \in \Omega$.
- (iii) For all $\omega \in \Omega$, $\mathbb{P}_{\mathcal{H}_{\mathbb{U}}}^{\text{do}(\mathbb{W}, \delta_\omega)}(\gamma; A) = K_{\mathbb{W}}(\omega; A)$ for $\mathbb{P}^{\text{do}(\mathbb{W}, \delta_\omega)}$ -almost all $\gamma \in \Omega$, by Proposition 3.3.23.

By the transitivity of $\mathbb{P}^{\text{do}(\mathbb{W}, \delta_\omega)}$ -almost sure equalities, the result follows. $\square$

Statement (ii) is implied by, but is weaker than, the absence of an active causal effect of $K_{\mathbb{U}}$ on A relative to $K_{\mathbb{W}}$ (Definition 3.3.15), in the same way that Theorem 3.4.3(b) was almost, but not quite, the absence of an active causal effect.

Taking $\mathbb{W} = \emptyset$ recovers Theorem 3.4.3: by the trivial intervention axiom (Axiom (i)) we have $\mathbb{P}^{\text{do}(\emptyset, \delta_\omega)} = \mathbb{P}$, and Lemma 3.2.4(i) gives $K_{\mathbb{U}}^{\text{do}(\emptyset, \delta_\omega)}(\cdot; A) = K_{\mathbb{U}}(\cdot; A)$, so the three statements become $K_{\mathbb{U}}(\cdot; A) = \mathbb{P}_{\mathcal{H}_{\mathbb{U}}}(\cdot, A)$, $K_{\mathbb{U}}(\cdot; A) = \mathbb{P}(A)$ and $\mathbb{P}_{\mathcal{H}_{\mathbb{U}}}(\cdot, A) = \mathbb{P}(A)$, all $\mathbb{P}$ -almost surely.

Just as we did in Definitions 3.3.21 and 3.3.22, let us combine the definitions of relative sources (Definition 3.4.9) and conditional sources (Definition 3.4.5).

Definition 3.4.12 (Conditional relative sources). Let $\mathbb{U}, \mathbb{W} \in \mathcal{P}(\mathbb{T})$, $A \in \mathcal{H}$ and $\mathcal{G} \subseteq \mathcal{H}$. We say that the causal kernel $K_{\mathbb{U}}$ is a *conditional* source of A given $\mathcal{G}$ *relative to* $K_{\mathbb{W}}$ if, for all $\omega \in \Omega$, there exists a source version $\mathbb{P}_{K_{\mathbb{U}}|\mathcal{G}}^{\text{do}(\mathbb{W}, \delta_\omega)}(\cdot; A)$ of $\mathbb{P}_{\mathcal{H}_{\mathbb{U}} \vee \mathcal{G}}^{\text{do}(\mathbb{W}, \delta_\omega)}(\cdot, A)$ that satisfies, for all $B \in \mathcal{H}_{\mathbb{U}}$ and all $G \in \mathcal{G}$,

$$K_{\mathbb{U} \cup \mathbb{W}}(\omega; A \cap B \cap G) = \int K_{\mathbb{U} \cup \mathbb{W}}(\omega; d\gamma) \mathbf{1}_{B \cap G}(\gamma) \mathbb{P}_{K_{\mathbb{U}}|\mathcal{G}}^{\text{do}(\mathbb{W}, \delta_\omega)}(\gamma; A).$$

We say that $K_{\mathbb{U}}$ is a conditional source of a sub- σ -algebra $\mathcal{F} \subseteq \mathcal{H}$ given $\mathcal{G}$ relative to $K_{\mathbb{W}}$ if it is a conditional source of all $A \in \mathcal{F}$ given $\mathcal{G}$ relative to $K_{\mathbb{W}}$.

Note that, as conditioning is involved, we took the route of defining conditional relative sources via the existence of source versions again, as in Definition 3.4.5. The intuition is as expected: once we intervene on $\mathcal{H}_{\mathbb{W}}$ and condition on $\mathcal{G}$, intervening further on $\mathcal{H}_{\mathbb{U}}$ is the same as conditioning on $\mathcal{H}_{\mathbb{U}}$. Indeed, Definition 3.4.12 says that, if $K_{\mathbb{U}}$ is a conditional source of A given $\mathcal{G}$ relative to $K_{\mathbb{W}}$, then a source version $\mathbb{P}_{K_{\mathbb{U}}|\mathcal{G}}^{\text{do}(\mathbb{W}, \delta_\omega)}(\cdot; A)$ is a version of the conditional probability $\mathbb{P}_{\mathcal{H}_{\mathbb{U}} \vee \mathcal{G}}^{\text{do}(\mathbb{U} \cup \mathbb{W}, \delta_\omega)}(\cdot, A)(\cdot, A) = \mathbb{P}_{\mathcal{G}}^{\text{do}(\mathbb{U} \cup \mathbb{W}, \delta_\omega)}(\cdot, A)$.

If $\mathbb{W} = \emptyset$, then we immediately recover conditional sources (Definition 3.4.5). If $\mathcal{G} = \{\emptyset, \Omega\}$, then $\mathbb{P}_{K_{\mathbb{U}}|\mathcal{G}}^{\text{do}(\mathbb{W}, \delta_\omega)}(\cdot; A) = \mathbb{P}_{K_{\mathbb{U}}}^{\text{do}(\mathbb{W}, \delta_\omega)}(\cdot; A)$ and $\mathbf{1}_B(\cdot)$ are $\mathcal{H}_{\mathbb{U}}$ -measurable, so by Theorem 3.1.2(c), the displayed equation in Definition 3.4.12 says, for all $\omega, \gamma \in \Omega$,

$$K_{\mathbb{U} \cup \mathbb{W}}((\gamma_{\mathbb{U}}, \omega_{\mathbb{W} \setminus \mathbb{U}}); A) = \mathbb{P}_{K_{\mathbb{U}}}^{\text{do}(\mathbb{W}, \delta_\omega)}(\gamma; A).$$

Here, the right-hand side only depends on ω through the $\mathbb{W}$ -component $\omega_{\mathbb{W}}$. Hence, for each value of $\omega_{\mathbb{U}}$, we can range through $\gamma \in \Omega$ to see that $\mathbb{P}_{K_{\mathbb{U}}}^{\text{do}(\mathbb{W}, \delta_\omega)}(\gamma; A)$ is uniquely defined for all $\gamma \in \Omega$. Using the fact that it is a version of $\mathbb{P}_{\mathcal{H}_{\mathbb{U}}}^{\text{do}(\mathbb{W}, \delta_\omega)}(\cdot, A)$, we have, for

all $B \in \mathcal{H}_{\mathbb{U}}$,

$$K_{\mathbb{W}}(\omega; A \cap B) = \int K_{\mathbb{W}}(\omega; d\gamma) \mathbf{1}_B(\gamma) K_{\mathbb{U} \cup \mathbb{W}}((\gamma_{\mathbb{U}}, \omega_{\mathbb{W} \setminus \mathbb{U}}); A),$$

recovering the definition of relative sources (Definition 3.4.9).

We end the chapter by showcasing the significance of the notions of conditional relative causal effects (Definition 3.3.21), conditional relative sources (Definition 3.4.12) and conditional causal independence (Definition 3.3.22). Let us take $\mathbb{U}, \mathbb{V}, \mathbb{W} \in \mathcal{P}(\mathbb{T})$, and let us fix $\omega \in \Omega$.

- If $K_{\mathbb{U}}$ does not have an active conditional causal effect on A conditioned on $\mathcal{H}_{\mathbb{V}}$ relative to $K_{\mathbb{W}}$, then from Definition 3.3.21 we know that $\mathbb{P}^{\text{do}(\mathbb{W}, \delta_{\omega})}$ and $\mathbb{P}^{\text{do}(\mathbb{U} \cup \mathbb{W}, \delta_{\omega})}$ are mutually absolutely continuous on $\mathcal{H}_{\mathbb{V}}$, and

$$\mathbb{P}_{\mathcal{H}_{\mathbb{V}}}^{\text{do}(\mathbb{U} \cup \mathbb{W}, \delta_{\omega})}(\cdot, A) = \mathbb{P}_{\mathcal{H}_{\mathbb{V}}}^{\text{do}(\mathbb{W}, \delta_{\omega})}(\cdot, A)$$

$\mathbb{P}^{\text{do}(\mathbb{W}, \delta_{\omega})}$ -almost surely (or equivalently, $\mathbb{P}^{\text{do}(\mathbb{U} \cup \mathbb{W}, \delta_{\omega})}$ -almost surely).

- If $K_{\mathbb{U}}$ is a conditional source of A given $\mathcal{H}_{\mathbb{V}}$ relative to $K_{\mathbb{W}}$, then by Definition 3.4.12, there exists a source version $\mathbb{P}_{K_{\mathbb{U}}|\mathcal{H}_{\mathbb{V}}}^{\text{do}(\mathbb{W}, \delta_{\omega})}(\cdot, A)$ of $\mathbb{P}_{\mathcal{H}_{\mathbb{U}} \vee \mathcal{H}_{\mathbb{V}}}^{\text{do}(\mathbb{W}, \delta_{\omega})}(\cdot, A)$ that, by Lemma 2.1.3 and Proposition 2.7.3(ii) applied to the defining identity, is also a version of $\mathbb{P}_{\mathcal{H}_{\mathbb{U}} \vee \mathcal{H}_{\mathbb{V}}}^{\text{do}(\mathbb{U} \cup \mathbb{W}, \delta_{\omega})}(\cdot, A)$; the latter in turn agrees $\mathbb{P}^{\text{do}(\mathbb{U} \cup \mathbb{W}, \delta_{\omega})}$ -almost surely with $\mathbb{P}_{\mathcal{H}_{\mathbb{V}}}^{\text{do}(\mathbb{U} \cup \mathbb{W}, \delta_{\omega})}(\cdot, A)$, by two applications of Lemma 3.1.4(i). We therefore write

$$\mathbb{P}_{\mathcal{H}_{\mathbb{V}}}^{\text{do}(\mathbb{U} \cup \mathbb{W}, \delta_{\omega})}(\cdot, A) = \mathbb{P}_{\mathcal{H}_{\mathbb{U}} \vee \mathcal{H}_{\mathbb{V}}}^{\text{do}(\mathbb{W}, \delta_{\omega})}(\cdot, A),$$

understood as the statement that a single function $\mathbb{P}_{K_{\mathbb{U}}|\mathcal{H}_{\mathbb{V}}}^{\text{do}(\mathbb{W}, \delta_{\omega})}(\cdot, A)$ is a version of the left-hand side $\mathbb{P}^{\text{do}(\mathbb{U} \cup \mathbb{W}, \delta_{\omega})}$ -almost surely and of the right-hand side $\mathbb{P}^{\text{do}(\mathbb{W}, \delta_{\omega})}$ -almost surely. Unlike the other two displays, this one cannot be read as an almost sure identity under either measure alone.

- If A and $\mathcal{H}_{\mathbb{U}}$ are causally independent under $K_{\mathbb{W}}$ conditioned on $\mathcal{H}_{\mathbb{V}}$, then Proposition 3.3.23(ii), applied with the causal kernel $K_{\mathbb{W}}$, the σ -algebras $\mathcal{F}_1 = \mathcal{H}_{\mathbb{U}}$ and $\mathcal{F}_2 = \sigma(A)$, and the conditioning σ -algebra $\mathcal{G} = \mathcal{H}_{\mathbb{V}}$, gives

$$\mathbb{P}_{\mathcal{H}_{\mathbb{U}} \vee \mathcal{H}_{\mathbb{V}}}^{\text{do}(\mathbb{W}, \delta_{\omega})}(\cdot, A) = \mathbb{P}_{\mathcal{H}_{\mathbb{V}}}^{\text{do}(\mathbb{W}, \delta_{\omega})}(\cdot, A)$$

$\mathbb{P}^{\text{do}(\mathbb{W}, \delta_{\omega})}$ -almost surely.

4 Causal spaces

In this chapter, we introduce *causal spaces* in full generality. Just like homogeneous causal spaces introduced in the previous chapter, a general causal space (Definition 4.1.1) is a mathematical structure that can carry *interventions*, the key notion behind causality. The causation encoded by a general causal space may be *heterogeneous*: the effect of an intervention can vary across subpopulations, depending on the pre-intervention information. For instance, the same medicine may have different effects depending on the patient's age. We formalise the notion of heterogeneity of causation in Section 4.2.

Subsequently, in Section 4.3, we introduce preservation of events, namely, the notion of an event staying at the pre-intervention outcome. In Section 4.4, we extend our theory to *targeted* causation, with the introduction of targeted causal kernels that can carry interventions on different parts of the system depending on the pre-intervention information.

4.1 The axioms

Just like Chapter 3, our starting point is the product measurable space $(\Omega, \mathcal{H}) = \otimes_{t \in \mathbb{T}} (\Omega_t, \mathcal{E}_t)$. In this chapter, to deal with the complications that arise from taking causal heterogeneity and targeting into account, we need some further technical preliminaries for measurability bookkeeping. First, we define a σ -algebra $\mathcal{T}$ on $\mathcal{P}(\mathbb{T})$ as follows. For each $t \in \mathbb{T}$, define the membership indicator function

$$\chi_t : \mathcal{P}(\mathbb{T}) \rightarrow \{0, 1\}, \quad \chi_t(\mathbb{S}) = \begin{cases} 0 & \text{if } t \notin \mathbb{S} \\ 1 & \text{if } t \in \mathbb{S}. \end{cases}$$

Then we define $\mathcal{T}$ as the smallest σ -algebra on $\mathcal{P}(\mathbb{T})$ such that χ_t is measurable for all $t \in \mathbb{T}$. Formally:

$$\mathcal{T} = \sigma(\chi_t : t \in \mathbb{T}) = \sigma(\{\{\mathbb{S} \in \mathcal{P}(\mathbb{T}) : t \in \mathbb{S}\} : t \in \mathbb{T}\}). \quad (4.1)$$

If $\mathbb{T}$ is finite, then $\mathcal{T} = \mathcal{P}(\mathcal{P}(\mathbb{T}))$, i.e. the set of all possible collections of subsets of $\mathbb{T}$. If $\mathbb{T}$ is countable, all singletons $\{\mathbb{S}\}$ are in $\mathcal{T}$, since each $\{\mathbb{S}\}$ is a countable intersection of the generating sets $\{\mathbb{S} \in \mathcal{P}(\mathbb{T}) : t \in \mathbb{S}\}$ and their complements, but in general, $\mathcal{T}$ is a strict subset of $\mathcal{P}(\mathcal{P}(\mathbb{T}))$. For uncountable $\mathbb{T}$, $\mathcal{T}$ is well-defined, but is far smaller than $\mathcal{P}(\mathcal{P}(\mathbb{T}))$, and does not in general contain all the singletons.

If a map $U : \Omega \rightarrow \mathcal{P}(\mathbb{T})$ is $\mathcal{H}/\mathcal{T}$ -measurable, it means that, for each $t \in \mathbb{T}$, the event $\{\eta \in \Omega : t \in U(\eta)\}$ is a measurable set in $\mathcal{H}$. Equivalently, U is $\mathcal{H}/\mathcal{T}$ -measurable if and only if all coordinate maps $\eta \mapsto \chi_t(U(\eta))$ are $\mathcal{H}$ -measurable. For uncountable $\mathbb{T}$,




Figure 4.1: The causal kernel $K_U((\eta, \omega); \cdot)$ (Definition 4.1.1) is the measure when the $\mathbb{U}$ -coordinates are fixed to the values ω_U , for the pre-intervention outcome η . The measure on $(\Omega, \mathcal{H}_U)$ collapses to the delta measure at ω_U regardless of the pre-intervention outcome η , consistent with the interventional determinism axiom (Axiom (ii)).

since the singletons are not in general contained in $\mathcal{T}$, the pre-images $U^{-1}(\{\mathbb{S}\})$ for each $\mathbb{S} \in \mathcal{P}(\mathbb{T})$ need not be measurable in general. We will make heavy use of this map U when considering *targeting*, in Sections 4.4 and 5.4.

Lastly, we define a σ -algebra $\mathcal{U}$ on $\mathcal{P}(\mathbb{T}) \times \Omega \times \Omega$ by

$$\mathcal{U} = (\mathcal{T} \otimes \mathcal{H} \otimes \{\emptyset, \Omega\}) \vee \sigma(\{(\mathbb{S}, \eta, \omega) : t \in \mathbb{S}, \omega \in H\} : t \in \mathbb{T}, H \in \mathcal{H}_t) \quad (4.2)$$

Semantically, $\mathcal{U}$ contains full information on the first two components $(\mathbb{S}, \eta)$, but only the ω -information on the selected coordinates $\mathbb{S}$. This σ -algebra $\mathcal{U}$, along with its many sub- σ -algebras and associated sections, will play a crucial role in ensuring clean measurability conditions throughout.

We are finally ready to introduce the formal definition of a causal space in full generality.

Definition 4.1.1 (Causal space). A *causal space* $(\Omega, \mathcal{H}, \mathbb{P}, \mathbb{K})$ consists of

- a product measurable space $(\Omega, \mathcal{H}) = \otimes_{t \in \mathbb{T}} (\Omega_t, \mathcal{E}_t)$;
- a probability measure $\mathbb{P}$ on $(\Omega, \mathcal{H})$; and
- a *causal mechanism* $\mathbb{K} = \{K_{\mathbb{S}} : \mathbb{S} \in \mathcal{P}(\mathbb{T})\}$, which is a collection of transition probability kernels $K_{\mathbb{S}}$, called *causal kernels*, from $(\Omega, \mathcal{H}) \otimes (\Omega, \mathcal{H}_{\mathbb{S}})$ into $(\Omega, \mathcal{H})$,

such that, for each $A \in \mathcal{H}$, the map $(\mathbb{S}, \eta, \omega) \mapsto K_{\mathbb{S}}((\eta, \omega); A) : \mathcal{P}(\mathbb{T}) \times \Omega \times \Omega \rightarrow [0, 1]$ is $\mathcal{U}$ -measurable.

The causal kernels $K_{\mathbb{S}}$ are required to satisfy the following axioms:

- (i) *trivial intervention*: for all $A \in \mathcal{H}$ and $\eta, \omega \in \Omega$,

$$K_{\emptyset}((\eta, \omega); A) = \mathbf{1}_A(\eta) = \delta_{\eta}(A);$$

- (ii) *interventional determinism*: for all $\mathbb{S} \in \mathcal{P}(\mathbb{T})$, $B \in \mathcal{H}_{\mathbb{S}}$ and $\eta, \omega \in \Omega$,

$$K_{\mathbb{S}}((\eta, \omega); B) = \mathbf{1}_B(\omega) = \delta_{\omega}(B).$$

We make a quick remark on the notation and the measurability condition. For fixed $\mathbb{S} \in \mathcal{P}(\mathbb{T})$, $\eta \in \Omega$ and $A \in \mathcal{H}$, the map $\omega \mapsto K_{\mathbb{S}}((\eta, \omega); A)$ is $\mathcal{H}_{\mathbb{S}}$ -measurable, and hence only dependent on the $\omega_{\mathbb{S}}$ component of ω . With an abuse of notation, we write $K_{\mathbb{S}}((\eta, \omega); A)$ and $K_{\mathbb{S}}((\eta, \omega_{\mathbb{S}}); A)$ interchangeably. The $\mathcal{U}$ -measurability condition of the joint map $(\mathbb{S}, \eta, \omega) \mapsto K_{\mathbb{S}}((\eta, \omega); A)$ is equivalent to saying that $\mathbb{K}$ is a transition probability kernel from $(\mathcal{P}(\mathbb{T}) \times \Omega \times \Omega, \mathcal{U})$ into $(\Omega, \mathcal{H})$. The definition of $\mathcal{U}$ in fact ensures that, for each fixed $\mathbb{S} \in \mathcal{P}(\mathbb{T})$ and $A \in \mathcal{H}$, the map $(\eta, \omega) \mapsto K_{\mathbb{S}}((\eta, \omega); A)$ is $(\mathcal{H} \otimes \mathcal{H}_{\mathbb{S}})$ -measurable.

The effect of fixing the $\mathbb{U}$ -coordinates to $\omega_{\mathbb{U}}$ is different, or *heterogeneous*, depending on the pre-intervention values, captured by the η argument (see Fig. 4.1 for an illustration). Heterogeneity is truly what sets general causal spaces apart from homogeneous ones introduced in Chapter 3, and will be discussed further in Section 4.2. In Fig. 4.1, each pre-intervention outcome value $\eta \in \Omega$ is visually illustrated as a person, but as discussed at the beginning of Chapter 3, each $\eta \in \Omega$ should not be viewed literally as a single person from the population, but rather, as a representative character with those attributes captured by η .

The formal definition of an intervention will be introduced in Chapter 5; for now, one can think of the causal kernels as encoding a particular type of intervention where one fixes the $\mathbb{U}$ -coordinates to a single value, $\omega_{\mathbb{U}}$. For the rest of this chapter, any mention of “interventions” should be understood in this sense.

The trivial intervention axiom (Axiom (i)) has the same semantic interpretation as the corresponding axiom in homogeneous causal spaces: that intervening on $\mathcal{H}_{\emptyset}$ changes nothing. In homogeneous causal spaces, where everything is stated at the population level, this meant that the probability measure on any event $A \in \mathcal{H}$ stayed at the observational measure $\mathbb{P}(A)$. In causal spaces as introduced in Definition 4.1.1, causal kernels encode information at the level of individual pre-intervention outcomes η , and hence saying that “nothing changes” means that any event $A \in \mathcal{H}$ really stays at this value η : $K_{\emptyset}((\eta, \omega); A) = \delta_{\eta}(A)$.

The interventional determinism axiom (Axiom (ii)) also has the same semantic interpretation as Axiom (ii), but with an extra implication for the heterogeneous argument η . Once the $\mathbb{S}$ -coordinates are given the values ω , then *at those coordinates*, the pre-intervention values η are completely overwritten.

The following is an analogue of Theorem 3.1.2 that characterises interventional determinism in several different forms, and has almost exactly the same proof, but we write it out for completeness.

Theorem 4.1.2 (Interventional determinism). *The interventional determinism axiom (Axiom (ii)) is equivalent to the following three statements.*

(a) For all $A \in \mathcal{H}$, all $B \in \mathcal{H}_{\mathbb{S}}$ and all $\eta, \omega \in \Omega$,

$$K_{\mathbb{S}}((\eta, \omega); A \cap B) = \mathbf{1}_B(\omega) K_{\mathbb{S}}((\eta, \omega); A).$$

(b) If $f : \Omega \rightarrow \overline{\mathbb{R}}_+$ is $\mathcal{H}$ -measurable, or if $f : \Omega \rightarrow \overline{\mathbb{R}}$ is $\mathcal{H}$ -measurable and integrable, then for all $\eta, \omega \in \Omega$,

$$\int K_{\mathbb{S}}((\eta, \omega); d\omega') f(\omega') = \int K_{\mathbb{S}}((\eta, \omega); d\omega') f((\omega_{\mathbb{S}}, \omega'_{\mathbb{T} \setminus \mathbb{S}})).$$

(c) Suppose that $g, h : \Omega \rightarrow \overline{\mathbb{R}}$ are $\mathcal{H}_{\mathbb{S}}$ - and $\mathcal{H}$ -measurable respectively. Suppose further that the following two conditions are satisfied.

- Either $gh : \Omega \rightarrow \overline{\mathbb{R}}_+$ or $gh : \Omega \rightarrow \overline{\mathbb{R}}$ is integrable.
- Either $h : \Omega \rightarrow \overline{\mathbb{R}}_+$ or $h : \Omega \rightarrow \overline{\mathbb{R}}$ is integrable.

Then for all $\eta, \omega \in \Omega$,

$$\int K_{\mathbb{S}}((\eta, \omega); d\omega') g(\omega') h(\omega') = g(\omega) \int K_{\mathbb{S}}((\eta, \omega); d\omega') h(\omega').$$

Proof. **Axiom (ii) $\implies$ (a)** We look at the following two cases separately.

- $\omega \notin B$: We have $\mathbf{1}_B(\omega) = 0$; see that, by Axiom (ii),

$$K_{\mathbb{S}}((\eta, \omega); A \cap B) \leq K_{\mathbb{S}}((\eta, \omega); B) = \mathbf{1}_B(\omega) = 0,$$

giving us $K_{\mathbb{S}}((\eta, \omega); A \cap B) = 0 = \mathbf{1}_B(\omega) K_{\mathbb{S}}((\eta, \omega); A)$.

- $\omega \in B$: We have $K_{\mathbb{S}}((\eta, \omega); B) = \mathbf{1}_B(\omega) = 1$ by Axiom (ii), i.e. B has probability 1 under the measure $K_{\mathbb{S}}((\eta, \omega); \cdot)$. Hence, for all events $A \in \mathcal{H}$,

$$K_{\mathbb{S}}((\eta, \omega); A \cap B) = K_{\mathbb{S}}((\eta, \omega); A) = \mathbf{1}_B(\omega) K_{\mathbb{S}}((\eta, \omega); A),$$

as required.

(a) $\implies$ (b) For each $\omega \in \Omega$, the map $\iota_{\omega_{\mathbb{S}}} \circ \pi_{\mathbb{T} \setminus \mathbb{S}} : \Omega \rightarrow \Omega$ overwrites the $\mathbb{S}$ -components: $\iota_{\omega_{\mathbb{S}}} \circ \pi_{\mathbb{T} \setminus \mathbb{S}}(\omega') = (\omega_{\mathbb{S}}, \omega'_{\mathbb{T} \setminus \mathbb{S}})$. As a composition of measurable functions, it is measurable. Now, for any measurable rectangle $B \cap C$ with $B \in \mathcal{H}_{\mathbb{S}}$ and $C \in \mathcal{H}_{\mathbb{T} \setminus \mathbb{S}}$, part (a) gives

$$K_{\mathbb{S}}((\eta, \omega); B \cap C) = \mathbf{1}_B(\omega) K_{\mathbb{S}}((\eta, \omega); C).$$

On the other hand, see that

$$(\iota_{\omega_S} \circ \pi_{T \setminus S})^{-1}(B \cap C) = \{\omega' \in \Omega : (\omega_S, \omega'_{T \setminus S}) \in B \cap C\} = \begin{cases} C & \text{if } \omega \in B \\ \emptyset & \text{otherwise,} \end{cases}$$

hence we also have

$$\begin{aligned} K_S((\eta, \omega); (\iota_{\omega_S} \circ \pi_{T \setminus S})^{-1}(B \cap C)) &= \begin{cases} K_S((\eta, \omega); C) & \text{if } \omega \in B \\ 0 & \text{otherwise} \end{cases} \\ &= \mathbf{1}_B(\omega) K_S((\eta, \omega); C), \end{aligned}$$

which is the same as $K_S((\eta, \omega); B \cap C)$ as derived from part (a) above. So the two measures $K_S((\eta, \omega); \cdot)$ and $K_S((\eta, \omega); (\iota_{\omega_S} \circ \pi_{T \setminus S})^{-1}(\cdot))$ coincide on the generating π -system of rectangles, and by the unique extension to all of $\mathcal{H} = \mathcal{H}_S \vee \mathcal{H}_{T \setminus S}$ given by Lemma 2.3.2, these two measures coincide on all of $\mathcal{H}$.

Now, take any $\mathcal{H}$ -measurable function $f : \Omega \rightarrow \bar{\mathbb{R}}_+$. Then by the above argument,

$$\begin{aligned} \int K_S((\eta, \omega); d\omega') f(\omega') &= \int K_S((\eta, \omega); (\iota_{\omega_S} \circ \pi_{T \setminus S})^{-1}(d\omega')) f(\omega') \\ &= \int K_S((\eta, \omega); d\omega') f(\iota_{\omega_S} \circ \pi_{T \setminus S}(\omega')) \\ &= \int K_S((\eta, \omega); d\omega') f((\omega_S, \omega'_{T \setminus S})), \end{aligned}$$

which is precisely (b).

(b) $\implies$ (c) In (b), let $f = gh$. Then

$$\begin{aligned} \int K_S((\eta, \omega); d\omega') g(\omega') h(\omega') &= \int K_S((\eta, \omega); d\omega') g((\omega_S, \omega'_{T \setminus S})) h((\omega_S, \omega'_{T \setminus S})) \\ &= g(\omega) \int K_S((\eta, \omega); d\omega') h((\omega_S, \omega'_{T \setminus S})) \quad \text{since } g \text{ is } \mathcal{H}_S\text{-measurable} \\ &= g(\omega) \int K_S((\eta, \omega); d\omega') h(\omega'), \end{aligned}$$

where we applied part (b) again on the last line with $f = h$. This is precisely (c).

(c) $\implies$ Axiom (ii) For any $B \in \mathcal{H}_S$, let $g = \mathbf{1}_B$ and $h = 1$. Then we have

$$K_S((\eta, \omega); B) = \int K_S((\eta, \omega); d\omega') \mathbf{1}_B(\omega') = \mathbf{1}_B(\omega),$$

as required. □

This theorem is a key result, and will be used again and again throughout the monograph.

In Remark 3.1.3, we highlighted the distinction between *almost sure* statements that are prevalent in probability theory, and statements that hold *for all* elements in Ω . The remark carries over to the heterogeneous kernels: each $K_{\mathbb{S}}((\eta, \omega); \cdot)$ is defined *for all* $\eta, \omega \in \Omega$. After interventions take place (Chapter 5), the sets which previously had zero measure may gain positive measure, which means that it is crucial for the causal kernels to be defined everywhere, to accommodate subsequent interventions.

4.2 Heterogeneity of causation

The heterogeneity of a causal kernel, encoded by its dependence on the pre-intervention argument η , may be measurable with respect to different bodies of pre-intervention information. For example, the effect of a medicine on blood pressure may depend on the age of the patient. Formally, we have the following definition.

Definition 4.2.1 (Heterogeneity family). For each event $A \in \mathcal{H}$, a σ -algebra $\mathcal{J} \subseteq \mathcal{H}$ is a *heterogeneity σ -algebra* of the causal kernel $K_{\mathbb{U}}$ on A if the map

$$(\eta, \omega) \mapsto K_{\mathbb{U}}((\eta, \omega); A)$$

is $(\mathcal{J} \otimes \mathcal{H}_{\mathbb{U}})$ -measurable. We then define the *heterogeneity family* of $K_{\mathbb{U}}$ on A as the collection of all of its heterogeneity σ -algebras on A , and denote it by $\Sigma_{\mathbb{U}}(A)$.

For a sub- σ -algebra $\mathcal{F} \subseteq \mathcal{H}$, a σ -algebra $\mathcal{J} \subseteq \mathcal{H}$ is a heterogeneity σ -algebra of $K_{\mathbb{U}}$ on $\mathcal{F}$ if $\mathcal{J} \in \Sigma_{\mathbb{U}}(A)$ for every $A \in \mathcal{F}$. We define the *heterogeneity family* of $K_{\mathbb{U}}$ on $\mathcal{F}$ as

$$\Sigma_{\mathbb{U}}(\mathcal{F}) := \bigcap_{A \in \mathcal{F}} \Sigma_{\mathbb{U}}(A);$$

in other words, it is the collection of all of its heterogeneity σ -algebras on $\mathcal{F}$.

If $\mathcal{J}$ is a heterogeneity σ -algebra of $K_{\mathbb{U}}$ on $\mathcal{H}$, we simply say that it is a heterogeneity σ -algebra of $K_{\mathbb{U}}$. We also simply write $\Sigma_{\mathbb{U}} = \Sigma_{\mathbb{U}}(\mathcal{H})$.

We say that $K_{\mathbb{U}}$ is *homogeneous* on A if $\{\emptyset, \Omega\} \in \Sigma_{\mathbb{U}}(A)$, and *heterogeneous* otherwise. Likewise, we say that $K_{\mathbb{U}}$ is homogeneous on $\mathcal{F}$ if $\{\emptyset, \Omega\} \in \Sigma_{\mathbb{U}}(\mathcal{F})$, and we simply say that $K_{\mathbb{U}}$ is homogeneous if $\{\emptyset, \Omega\} \in \Sigma_{\mathbb{U}}$.

Note that $\mathcal{H}$ is a heterogeneity σ -algebra of all causal kernels on all events, since, by the definition of causal kernels (Definition 4.1.1), the map $(\eta, \omega) \mapsto K_{\mathbb{U}}((\eta, \omega); A)$ is $(\mathcal{H} \otimes \mathcal{H}_{\mathbb{U}})$ -measurable. Hence, a heterogeneity family is never empty. Moreover, $\Sigma_{\mathbb{U}}(A)$ is upward closed: if $\mathcal{J} \in \Sigma_{\mathbb{U}}(A)$ and $\mathcal{J} \subseteq \mathcal{J}' \subseteq \mathcal{H}$, then $\mathcal{J}' \in \Sigma_{\mathbb{U}}(A)$.

Semantically, the heterogeneity family $\Sigma_{\mathbb{U}}(A)$ collects all bodies of pre-intervention information that suffice to determine how the effect on A of intervening on the $\mathbb{U}$ -coordinates varies. For example, if $\mathbb{U}$ is a coordinate representing a medicine, and its effect on A (say, blood pressure) depends only on the age of the patient, at least among the variables explicitly modelled in the space, then the σ -algebra generated by

age belongs to $\Sigma_{\mathbb{U}}(A)$; and, by upward closure, so does any larger σ -algebra. Thus, $\Sigma_{\mathbb{U}}(A)$ should be understood not as a single distinguished σ -algebra, but as the collection of all bodies of pre-intervention information that are sufficient to encode the dependence of the intervention effect on the pre-intervention state.

Remark 4.2.2. It would have been clean to define a single heterogeneity σ -algebra of $K_{\mathbb{U}}$ on A as the intersection of all σ -algebras $\mathcal{J}$ such that $(\eta, \omega) \mapsto K_{\mathbb{U}}((\eta, \omega); A)$ is $(\mathcal{J} \otimes \mathcal{H}_{\mathbb{U}})$ -measurable. It is true that, if $(\eta, \omega) \mapsto K_{\mathbb{U}}((\eta, \omega); A)$ is $(\mathcal{J} \otimes \mathcal{H}_{\mathbb{U}})$ -measurable for all $\mathcal{J} \in \Sigma_{\mathbb{U}}(A)$, then it is $\cap_{\mathcal{J} \in \Sigma_{\mathbb{U}}(A)} (\mathcal{J} \otimes \mathcal{H}_{\mathbb{U}})$ -measurable. Since an arbitrary intersection of σ -algebras is a σ -algebra (Lemma 2.1.2), we also have that the intersection $\cap_{\mathcal{J} \in \Sigma_{\mathbb{U}}(A)} \mathcal{J}$ is a σ -algebra. However, unless additional regularity assumptions are imposed (such as countable generation), it is not guaranteed that we have

$$(\cap_{\mathcal{J} \in \Sigma_{\mathbb{U}}(A)} \mathcal{J}) \otimes \mathcal{H}_{\mathbb{U}} = \cap_{\mathcal{J} \in \Sigma_{\mathbb{U}}(A)} (\mathcal{J} \otimes \mathcal{H}_{\mathbb{U}}).$$

In this monograph, we prefer to keep things as general as possible, and not to impose such structural assumptions. Hence, the heterogeneity family $\Sigma_{\mathbb{U}}(A)$ is the quantity we use to capture the heterogeneity of causation.

If $K_{\mathbb{U}}$ is homogeneous on A , then for each $\omega \in \Omega$, the map $\eta \mapsto K_{\mathbb{U}}((\eta, \omega); A)$ is constant, so that fixing the $\mathbb{U}$ -coordinates to the values $\omega_{\mathbb{U}}$ has the same effect on A uniformly across all pre-intervention states η , at least insofar as these are represented by the variables in the model.

In some situations, homogeneity can hold in the absolute sense. For (an admittedly extreme) example, suppose that, in the product measurable space, we have one coordinate represent a person taking some medicine, another representing the number that comes up when they roll a die, and some other coordinates representing other attributes of the person. Then clearly, there is no causation from the medicine coordinate to the dice roll, and whatever attributes that the person had before intervention do not affect the fact that any intervention on the medicine coordinate will leave the measure on the dice roll coordinate intact.

More pertinently, homogeneity of causal kernels will often hold in a marginal sense. In the medicine and blood pressure example, suppose that the causation from medicine to blood pressure varies only depending on the age of the person. But if the age is not included in the model (i.e. it is marginalised out), the heterogeneity that arises from age is absorbed in the measure given by the causal kernel, and the causation is homogeneous with respect to the coordinates included in the model.

Heterogeneity σ -algebras need not be a component of the product measurable space, i.e. of the form $\mathcal{H}_{\mathbb{S}}$ for some $\mathbb{S} \in \mathcal{P}(\mathbb{T})$. Also, a heterogeneity σ -algebra $\mathcal{J} \in \Sigma_{\mathbb{U}}(\mathcal{F})$ can intersect with $\mathcal{H}_{\mathbb{U}}$ or $\mathcal{F}$, meaning that the effect of the treatment can depend on the observational state of the treated coordinates and/or of the events on whom we are interested in the effect. For example, the blood pressure of a patient after taking a medicine can depend on what it was before, unless the medicine is powerful enough to completely ignore what it was before. Note that, in this case, the causation will not be homogeneous even if we only have medicine and blood pressure in the model, and all other variables are marginalised out.

Let us see what the axioms of causal spaces (Axioms (i) and (ii) in Definition 4.1.1) imply on the homogeneity/heterogeneity of the corresponding causal kernels.

- (i) The kernel $K_\emptyset((\eta, \omega); \cdot)$ is required to be $\delta_\eta(\cdot)$ for all $\eta, \omega \in \Omega$ by Axiom (i), and hence is heterogeneous on all non-trivial events $A \in \mathcal{H}$. For each event A , we have

$$A \in \mathcal{J} \quad \text{if and only if} \quad \mathcal{J} \in \Sigma_\emptyset(A),$$

and for a sub- σ -algebra $\mathcal{F} \subseteq \mathcal{H}$, we have

$$\mathcal{F} \subseteq \mathcal{J} \quad \text{if and only if} \quad \mathcal{J} \in \Sigma_\emptyset(\mathcal{F}).$$

Semantically, when one intervenes on the empty set (i.e. when one does not do anything), the pre-intervention values η remain at those values; hence, in order to know whether or not A occurs for η , it is necessary to have information about A in the pre-intervention state.

- (ii) On the other hand, the interventional determinism axiom (Axiom (ii)) forces the causal kernel $K_\mathbb{U}((\eta, \omega); \cdot)$ to be $\delta_\omega(\cdot)$ on events in $\mathcal{H}_\mathbb{U}$, which is independent of η . Hence, every causal kernel $K_\mathbb{U}$ is homogeneous on the corresponding σ -algebra $\mathcal{H}_\mathbb{U}$, i.e. $\{\emptyset, \Omega\} \in \Sigma_\mathbb{U}(\mathcal{H}_\mathbb{U})$. By upward closure, this implies that, in fact, every sub- σ -algebra of $\mathcal{H}$ belongs to $\Sigma_\mathbb{U}(\mathcal{H}_\mathbb{U})$. Semantically, once we fix the $\mathbb{U}$ -coordinates to a particular value $\omega_\mathbb{U}$, then we do not need to know any pre-intervention information to know that they will remain at $\omega_\mathbb{U}$ after intervention, thanks to Axiom (ii).

We can *homogenise* causal kernels $K_\mathbb{S}$ to force them to be homogeneous, by integrating out the η argument with respect to the observational measure $\mathbb{P}$.

Definition 4.2.3 (Homogenisation). For each $\mathbb{S} \in \mathcal{P}(\mathbb{T})$, we define the *homogenisation* of the causal kernel $K_\mathbb{S}$ as

$$\overline{K}_\mathbb{S}(\omega; A) = \int \mathbb{P}(d\eta) K_\mathbb{S}((\eta, \omega); A).$$

Given a causal space $(\Omega, \mathcal{H}, \mathbb{P}, \mathbb{K})$, we define the *homogenisation* $\overline{\mathbb{K}}$ of the causal mechanism $\mathbb{K}$ by homogenising all of its causal kernels:

$$\overline{\mathbb{K}} = \{\overline{K}_\mathbb{S} : \mathbb{S} \in \mathcal{P}(\mathbb{T})\}.$$

Homogenisation is not an intrinsic operation on $K_\mathbb{S}$ alone; it is a $\mathbb{P}$ -dependent population averaging of heterogeneous responses. The homogenised kernel $\overline{K}_\mathbb{S}$ no longer depends on the η argument, and hence is omitted from the notation. Semantically, the homogenised kernel aggregates the effect of fixing the $\mathbb{S}$ -coordinates to $\omega_\mathbb{S}$ over the population, by integrating out the η argument with respect to $\mathbb{P}$ —see Fig. 4.2. Clearly, if the kernel $K_\mathbb{S}$ was already homogeneous on A , then we have $\overline{K}_\mathbb{S}(\omega; A) = K_\mathbb{S}((\eta, \omega); A)$ for all $\eta, \omega \in \Omega$. Homogenised causal kernels will play an important role throughout the monograph.




Figure 4.2: The homogenised causal kernel $\bar{K}_{\mathbb{U}}(\omega; \cdot)$ (Definition 4.2.3) is the measure when intervening on the coordinates $\mathbb{U}$ by fixing them to the value $\omega_{\mathbb{U}}$, averaged over the pre-intervention outcome η with respect to the population measure $\mathbb{P}$: $\bar{K}_{\mathbb{U}}(\omega; \cdot) = \int \mathbb{P}(d\eta) K_{\mathbb{U}}((\eta, \omega); \cdot)$. The homogenised kernel still satisfies interventional determinism (Proposition 4.2.4), and hence the measure on $\mathcal{H}_{\mathbb{U}}$ is still the delta measure δ_{ω} .

The homogenised causal mechanism $\bar{\mathbb{K}}$ (Definition 4.2.3) obtained by homogenising all of its constituent causal kernels is a valid causal mechanism of a homogeneous causal space, as introduced in the warm-up chapter (Definition 3.1.1).

Proposition 4.2.4. *Given a causal space $(\Omega, \mathcal{H}, \mathbb{P}, \mathbb{K})$, its homogenisation $(\Omega, \mathcal{H}, \mathbb{P}, \bar{\mathbb{K}})$ is a valid homogeneous causal space (Definition 3.1.1).*

Proof. We first show that each $\bar{K}_{\mathbb{S}}$ is a transition probability kernel from $(\Omega, \mathcal{H}_{\mathbb{S}})$ into $(\Omega, \mathcal{H})$.

Fix $\omega \in \Omega$. Then $\bar{K}_{\mathbb{S}}(\omega; \Omega) = \int \mathbb{P}(d\eta) K_{\mathbb{S}}((\eta, \omega); \Omega) = 1$, and for pairwise disjoint $A_1, A_2, \dots \in \mathcal{H}$, the countable additivity of each measure $K_{\mathbb{S}}((\eta, \omega); \cdot)$ gives

$$\begin{aligned} \bar{K}_{\mathbb{S}}(\omega; \sqcup_{n \in \mathbb{N}} A_n) &= \int \mathbb{P}(d\eta) \sum_{n \in \mathbb{N}} K_{\mathbb{S}}((\eta, \omega); A_n) \\ &= \sum_{n \in \mathbb{N}} \int \mathbb{P}(d\eta) K_{\mathbb{S}}((\eta, \omega); A_n) \\ &= \sum_{n \in \mathbb{N}} \bar{K}_{\mathbb{S}}(\omega; A_n). \end{aligned}$$

Hence $A \mapsto \bar{K}_{\mathbb{S}}(\omega; A)$ is a probability measure on $(\Omega, \mathcal{H})$ for each $\omega \in \Omega$.

Now fix $A \in \mathcal{H}$. By the definition of causal kernels (Definition 4.1.1), the map $(\eta, \omega) \mapsto K_{\mathbb{S}}((\eta, \omega); A)$ is $(\mathcal{H} \otimes \mathcal{H}_{\mathbb{S}})$ -measurable, so by Lemma 2.5.5 the map

$$\omega \mapsto \int \mathbb{P}(d\eta) K_{\mathbb{S}}((\eta, \omega); A) = \overline{K}_{\mathbb{S}}(\omega; A)$$

is $\mathcal{H}_{\mathbb{S}}$ -measurable. Therefore $\overline{K}_{\mathbb{S}}$ is a transition probability kernel from $(\Omega, \mathcal{H}_{\mathbb{S}})$ into $(\Omega, \mathcal{H})$, as required by Definition 3.1.1.

Now we check the two axioms of causal kernels. By the trivial intervention axiom of $K_{\emptyset}$, we have $K_{\emptyset}((\eta, \omega); \cdot) = \delta_{\eta}(\cdot)$ for all $\eta, \omega \in \Omega$. Hence, for all $\omega \in \Omega$ and all $A \in \mathcal{H}$,

$$\overline{K}_{\emptyset}(\omega; A) = \int \mathbb{P}(d\eta) K_{\emptyset}((\eta, \omega); A) = \int \mathbb{P}(d\eta) \delta_{\eta}(A) = \int \mathbb{P}(d\eta) \mathbf{1}_A(\eta) = \mathbb{P}(A),$$

which is precisely Axiom (i) in Definition 3.1.1.

For the interventional determinism axiom, take any $\mathbb{S} \in \mathcal{P}(\mathbb{T})$, any $B \in \mathcal{H}_{\mathbb{S}}$ and any $\omega \in \Omega$. By Axiom (ii) of Definition 4.1.1, we have $K_{\mathbb{S}}((\eta, \omega); B) = \mathbf{1}_B(\omega)$ for every $\eta \in \Omega$, so that

$$\overline{K}_{\mathbb{S}}(\omega; B) = \int \mathbb{P}(d\eta) K_{\mathbb{S}}((\eta, \omega); B) = \int \mathbb{P}(d\eta) \mathbf{1}_B(\omega) = \mathbf{1}_B(\omega),$$

which is precisely Axiom (ii) of Definition 3.1.1. $\square$

Note that Proposition 4.2.4 reconciles the two forms taken by the trivial intervention axiom. In the general setting, Axiom (i) requires $K_{\emptyset}((\eta, \omega); \cdot) = \delta_{\eta}(\cdot)$: doing nothing leaves the pre-intervention outcome η exactly where it was. In the homogeneous setting of the warm-up chapter, Axiom (i) instead requires $K_{\emptyset}(\omega; A) = \mathbb{P}(A)$: doing nothing leaves the population measure $\mathbb{P}$ intact. The proof above shows that the second is the homogenisation of the first, the individual-level statement averaging to the population-level one. The apparent discrepancy between the two axioms is therefore precisely the loss of individual-level information that homogenisation effects.

Finally, the following consequence of the interventional determinism axiom is a useful one that will streamline proofs later in the monograph.

Lemma 4.2.5. *Suppose that $\mathcal{J} \in \Sigma_{\mathbb{U}}(A)$ is a heterogeneity σ -algebra of the causal kernel $K_{\mathbb{U}}$ on an event $A \in \mathcal{H}$. Let us take any $B \in \mathcal{H}_{\mathbb{U}}$. Then $\mathcal{J}$ is also a heterogeneity σ -algebra of $K_{\mathbb{U}}$ on $A \cap B$, i.e. $\mathcal{J} \in \Sigma_{\mathbb{U}}(A \cap B)$.*

Hence, $\Sigma_{\mathbb{U}}(A) \subseteq \Sigma_{\mathbb{U}}(A \cap B)$.

Proof. By the interventional determinism axiom (Theorem 4.1.2(a)), we have

$$K_{\mathbb{U}}((\eta, \omega); A \cap B) = \mathbf{1}_B(\omega) K_{\mathbb{U}}((\eta, \omega); A).$$

Here, the map $(\eta, \omega) \mapsto \mathbf{1}_B(\omega)$ is $(\{\emptyset, \Omega\} \otimes \mathcal{H}_{\mathbb{U}})$ -measurable, so trivially it is also $(\mathcal{J} \otimes \mathcal{H}_{\mathbb{U}})$ -measurable. By the definition of heterogeneity σ -algebras, the map $(\eta, \omega) \mapsto K_{\mathbb{U}}((\eta, \omega); A)$ is also $(\mathcal{J} \otimes \mathcal{H}_{\mathbb{U}})$ -measurable. Hence, the map

$$(\eta, \omega) \mapsto K_{\mathbb{U}}((\eta, \omega); A \cap B)$$

is $(\mathcal{J} \otimes \mathcal{H}_{\mathbb{U}})$ -measurable. So by definition, $\mathcal{J} \in \Sigma_{\mathbb{U}}(A \cap B)$. $\square$




Figure 4.3: Weak preservation (Definition 4.3.1 and Eq. (4.3)). If fixing the $\mathbb{U}$ -coordinates to any value $\omega_{\mathbb{U}}$ for any pre-intervention outcome $\eta \in \Omega$ preserves the measure of A at δ_η (i.e. $K_{\mathbb{U}}((\eta, \omega); A) = \delta_\eta(A)$ for all $\eta, \omega \in \Omega$), then $K_{\mathbb{U}}$ weakly preserves A .

4.3 Preservation

Let us return to the trivial intervention axiom (Axiom (i)). We can describe the behaviour of $K_\emptyset$ by saying that all events $A \in \mathcal{H}$ are *preserved* at the pre-intervention outcome η — A is certain if $\eta \in A$, and impossible otherwise. This is a property that may also hold for other kernels, and we give this as a definition.

Definition 4.3.1 (Preservation). We say that the causal kernel $K_{\mathbb{U}}$ *weakly preserves* an event $A \in \mathcal{H}$ if, for all $\eta, \omega \in \Omega$,

$$K_{\mathbb{U}}((\eta, \omega); A) = \delta_\eta(A) = \mathbf{1}_A(\eta). \quad (4.3)$$

We say that $K_{\mathbb{U}}$ *strongly preserves* A if, for all $\mathbb{S} \in \mathcal{P}(\mathbb{T})$ and all $\eta, \omega \in \Omega$,

$$K_{\mathbb{S}}((\eta, \omega); A) = K_{\mathbb{S} \setminus \mathbb{U}}((\eta, \omega); A). \quad (4.4)$$

For a sub- σ -algebra $\mathcal{F} \subseteq \mathcal{H}$, we say that $K_{\mathbb{U}}$ weakly/strongly preserves $\mathcal{F}$ if it weakly/strongly preserves every $A \in \mathcal{F}$.

Semantically, if $K_{\mathbb{U}}$ weakly preserves A , then *no matter what* values $\omega_{\mathbb{U}}$ we fix the $\mathbb{U}$ -coordinates to, the event A occurs if and only if the pre-intervention outcome η is in A . In other words, intervening on the $\mathbb{U}$ -coordinates has no effect on A —see Fig. 4.3 for an illustration. Strong preservation goes one step further— $K_{\mathbb{U}}$ strongly preserves A if intervening on the $\mathbb{S}$ -coordinates to ω is the same as doing so without the $\mathbb{U}$ -components.




Figure 4.4: Strong preservation (Definition 4.3.1 and Eq. (4.4)). If fixing any $\mathbb{S}$ -coordinates to any value $\omega_{\mathbb{S}}$ for any pre-intervention outcome $\eta \in \Omega$ is the same as fixing the $(\mathbb{S} \setminus \mathbb{U})$ -coordinates to the same (marginal) values $\omega_{\mathbb{S} \setminus \mathbb{U}}$ (i.e. $K_{\mathbb{U}}((\eta, \omega); A) = K_{\mathbb{S} \setminus \mathbb{U}}((\eta, \omega); A)$ for all $\mathbb{S} \in \mathcal{P}(\mathbb{T})$ and all $\eta, \omega \in \Omega$), then $K_{\mathbb{U}}$ strongly preserves A .

In other words, the $\mathbb{U}$ -coordinates cannot have any effect on A under interventions on any other coordinates—see Fig. 4.4 for an illustration.

Accordingly, weak preservation is a property of the single causal kernel $K_{\mathbb{U}}$, whereas strong preservation depends on the entire causal mechanism $\mathbb{K}$. These are precisely the heterogeneous, η -level analogues of the absence of active causal effects and no causal effects (Definition 3.3.1).

We can see that the trivial intervention axiom (Axiom (i)) is precisely the statement that $K_{\emptyset}$ weakly preserves every event in $\mathcal{H}$. It can also be seen immediately that $K_{\emptyset}$ strongly preserves all events $A \in \mathcal{H}$, and hence strongly preserves $\mathcal{H}$, since $\mathbb{S} \setminus \emptyset = \mathbb{S}$ for all $\mathbb{S} \in \mathcal{P}(\mathbb{T})$. On the other hand, for $\mathbb{U} \neq \emptyset$, $K_{\mathbb{U}}$ cannot preserve any non-trivial event in the corresponding σ -algebra $\mathcal{H}_{\mathbb{U}}$, due to the interventional determinism axiom (Axiom (ii)). Specifically, for $A \in \mathcal{H}_{\mathbb{U}}$, Axiom (ii) implies that $K_{\mathbb{U}}((\eta, \omega); A) = \mathbf{1}_A(\omega)$ for all $\eta, \omega \in \Omega$. If it also weakly preserved A , then we would also have $K_{\mathbb{U}}((\eta, \omega); A) = \mathbf{1}_A(\eta)$ for all $\eta, \omega \in \Omega$, meaning we would have to have $\mathbf{1}_A(\eta) = \mathbf{1}_A(\omega)$, which is only true if $A = \emptyset$ or $A = \Omega$. The notion of preservation will play a key role in many results throughout the monograph.

As the name suggests, strong preservation implies weak preservation, but not vice versa. This is the content of the next lemma, and is analogous to the result in Section 3.3 that no causal effect implies the absence of an active causal effect.

Lemma 4.3.2. *If the causal kernel $K_{\mathbb{U}}$ strongly preserves an event A , then it also weakly preserves A .*

Proof. Apply the definition of strong preservation (Eq. (4.4)) with $\mathbb{S} = \mathbb{U}$ and the trivial intervention axiom (Axiom (i)) to see that

$$\begin{aligned} K_{\mathbb{U}}((\eta, \omega); A) &= K_{\mathbb{U} \setminus \mathbb{U}}((\eta, \omega); A) \\ &= K_{\emptyset}((\eta, \omega); A) \\ &= \delta_{\eta}(A). \end{aligned}$$

Hence, $K_{\mathbb{U}}$ weakly preserves A . $\square$

If $K_{\mathbb{U}}$ (weakly) preserves an event A , then we have $K_{\mathbb{U}}((\eta, \omega); A) = \mathbf{1}_A(\eta)$ for all $\eta, \omega \in \Omega$, and hence $\{\emptyset, \Omega\} \notin \Sigma_{\mathbb{U}}(A)$, and so the kernel $K_{\mathbb{U}}$ is heterogeneous on A (in the sense of Definition 4.2.1), for all A that is neither $\emptyset$ nor Ω .

The intuition behind the case of weak preservation without strong preservation is the same as the intuition behind dormant causal effect, as introduced in Definition 3.3.1, but at the level of individual pre-intervention outcomes, rather than at the population level.

Weak preservation is already a very strong statement, that interventions on the $\mathbb{U}$ -coordinates cannot have any effect on A . For most purposes throughout the monograph, weak preservation is sufficient to derive the results we seek. In the next proposition, we show two equivalent forms of weak preservation. The first considers the causal kernel evaluated on intersection of events, one of which it weakly preserves; it is similar to the equivalent form of the interventional determinism axiom given in Theorem 4.1.2(a). The second is an integral form, using a standard monotone convergence theorem argument.

Proposition 4.3.3 (Weak preservation). *Let $\mathbb{U} \in \mathcal{P}(\mathbb{T})$ and $\mathcal{F} \subseteq \mathcal{H}$. The following three statements are equivalent.*

- (a) $K_{\mathbb{U}}$ weakly preserves $\mathcal{F}$.
- (b) For any pair of events $A \in \mathcal{F}$ and $D \in \mathcal{H}$,

$$K_{\mathbb{U}}((\eta, \omega); A \cap D) = \mathbf{1}_A(\eta) K_{\mathbb{U}}((\eta, \omega); D) \quad \text{for all } \eta, \omega \in \Omega.$$

- (c) For any $\mathcal{F}$ -measurable function $f : \Omega \rightarrow \overline{\mathbb{R}}_+$ and any $\mathcal{H}$ -measurable function $g : \Omega \rightarrow \overline{\mathbb{R}}_+$,

$$\int K_{\mathbb{U}}((\eta, \omega); d\omega') f(\omega') g(\omega') = f(\eta) \int K_{\mathbb{U}}((\eta, \omega); d\omega') g(\omega') \quad \text{for all } \eta, \omega \in \Omega.$$

Proof. **(a) $\implies$ (b)** We look at the following two cases separately.

- $\eta \notin A$: We have $\mathbf{1}_A(\eta) = 0$; since $K_{\mathbb{U}}$ weakly preserves $A \in \mathcal{F}$,

$$K_{\mathbb{U}}((\eta, \omega); A \cap D) \leq K_{\mathbb{U}}((\eta, \omega); A) = \mathbf{1}_A(\eta) = 0,$$

giving us $K_{\mathbb{U}}((\eta, \omega); A \cap D) = 0 = \mathbf{1}_A(\eta) K_{\mathbb{U}}((\eta, \omega); D)$.

- $\eta \in A$: We have $K_{\mathbb{U}}((\eta, \omega); A) = \mathbf{1}_A(\eta) = 1$ since $K_{\mathbb{U}}$ weakly preserves $A \in \mathcal{F}$, i.e. A has probability 1 under the measure $K_{\mathbb{U}}((\eta, \omega); \cdot)$. Hence, for all events $D \in \mathcal{H}$,

$$K_{\mathbb{U}}((\eta, \omega); A \cap D) = K_{\mathbb{U}}((\eta, \omega); D) = \mathbf{1}_A(\eta) K_{\mathbb{U}}((\eta, \omega); D),$$

as required.

- (b) $\implies$ (c)** Recall from Theorem 2.2.6 that a positive function is measurable if and only if it is the limit of an increasing sequence of positive simple functions. Since f and g are positive and $\mathcal{F}$ -measurable and $\mathcal{H}$ -measurable respectively, we can write them as an increasing limit of positive simple functions, say $f = \lim_{n \rightarrow \infty} f_n$ and $g = \lim_{n \rightarrow \infty} g_n$ for some $(f_n)_{n \in \mathbb{N}}$ and $(g_n)_{n \in \mathbb{N}}$ with

$$f_n = \sum_{i_n=1}^{m_n} a_{i_n} \mathbf{1}_{A_{i_n}}, \quad g_n = \sum_{j_n=1}^{\ell_n} d_{j_n} \mathbf{1}_{D_{j_n}}$$

where $a_{i_n}, d_{j_n} \in [0, \infty)$ and $A_{i_n} \in \mathcal{F}$, $D_{j_n} \in \mathcal{H}$. Then by the monotone convergence theorem (Theorem 2.5.1), for all $\eta, \omega \in \Omega$,

$$\begin{aligned} \int K_{\mathbb{U}}((\eta, \omega); d\omega') f(\omega') g(\omega') &= \int K_{\mathbb{U}}((\eta, \omega); d\omega') \left(\lim_{n \rightarrow \infty} f_n(\omega') g_n(\omega') \right) \\ &= \lim_{n \rightarrow \infty} \sum_{i_n=1}^{m_n} \sum_{j_n=1}^{\ell_n} a_{i_n} d_{j_n} \int K_{\mathbb{U}}((\eta, \omega); d\omega') \mathbf{1}_{A_{i_n}}(\omega') \mathbf{1}_{D_{j_n}}(\omega') \\ &= \lim_{n \rightarrow \infty} \sum_{i_n=1}^{m_n} \sum_{j_n=1}^{\ell_n} a_{i_n} d_{j_n} K_{\mathbb{U}}((\eta, \omega); A_{i_n} \cap D_{j_n}). \end{aligned}$$

By (b), we have, for all $\eta, \omega \in \Omega$,

$$\begin{aligned} \int K_{\mathbb{U}}((\eta, \omega); d\omega') f(\omega') g(\omega') &= \lim_{n \rightarrow \infty} \sum_{i_n=1}^{m_n} \sum_{j_n=1}^{\ell_n} a_{i_n} d_{j_n} \mathbf{1}_{A_{i_n}}(\eta) K_{\mathbb{U}}((\eta, \omega); D_{j_n}) \\ &= \left(\lim_{n \rightarrow \infty} f_n(\eta) \right) \left(\lim_{n \rightarrow \infty} \sum_{j_n=1}^{\ell_n} d_{j_n} K_{\mathbb{U}}((\eta, \omega); D_{j_n}) \right) \\ &= f(\eta) \left(\lim_{n \rightarrow \infty} \int K_{\mathbb{U}}((\eta, \omega); d\omega') \sum_{j_n=1}^{\ell_n} d_{j_n} \mathbf{1}_{D_{j_n}}(\omega') \right) \\ &= f(\eta) \int K_{\mathbb{U}}((\eta, \omega); d\omega') g(\omega'), \end{aligned}$$

where we applied the monotone convergence theorem again on the last line.

- (c) $\implies$ (a)** Let $f = \mathbf{1}_A$ for each $A \in \mathcal{F}$ and $g = 1$ to obtain the result immediately. $\square$




Figure 4.5: Weak relative preservation (Definition 4.3.5 and Eq. (4.5)). If fixing the $(\mathbb{U} \cup \mathbb{W})$ -coordinates to any value $\omega_{\mathbb{U}\mathbb{W}}$ for any pre-intervention outcome $\eta \in \Omega$ is the same as fixing the $\mathbb{W}$ -coordinates to $\omega_{\mathbb{W}}$ for the measure of A (i.e. $K_{\mathbb{U}\mathbb{W}}((\eta, \omega); A) = K_{\mathbb{W}}((\eta, \omega); A)$ for all $\eta, \omega \in \Omega$), then $K_{\mathbb{U}}$ weakly preserves A relative to $K_{\mathbb{W}}$.

It may seem natural that, if $K_{\mathbb{U}}$ preserves A , then for all subsets $\mathbb{V} \subseteq \mathbb{U}$, $K_{\mathbb{V}}$ should also preserve A . Unfortunately, since the causal kernels are each primitive objects of the space without any relationship between them a priori, this is not necessarily true for weak preservation. This monotonicity is, however, true for strong preservation, which does implicitly impose relationships among the causal kernels. We formalise this result in the following lemma.

Lemma 4.3.4. *If $K_{\mathbb{U}}$ strongly preserves A and $\mathbb{V} \subseteq \mathbb{U}$, then $K_{\mathbb{V}}$ strongly preserves A .*

Proof. Take any $\mathbb{S} \in \mathcal{P}(\mathbb{T})$. Applying Eq. (4.4) with $\mathbb{S}$ and $\mathbb{S} \setminus \mathbb{V}$ respectively, we have the following two identities for all $\eta, \omega \in \Omega$:

$$K_{\mathbb{S}}((\eta, \omega); A) = K_{\mathbb{S} \setminus \mathbb{U}}((\eta, \omega); A), \quad K_{\mathbb{S} \setminus \mathbb{V}}((\eta, \omega); A) = K_{(\mathbb{S} \setminus \mathbb{V}) \setminus \mathbb{U}}((\eta, \omega); A).$$

Since $\mathbb{V} \subseteq \mathbb{U}$, we have $(\mathbb{S} \setminus \mathbb{V}) \setminus \mathbb{U} = \mathbb{S} \setminus \mathbb{U}$, which means that the right-hand sides of both equations are the same. Hence, for all $\eta, \omega \in \Omega$, we must have

$$K_{\mathbb{S}}((\eta, \omega); A) = K_{\mathbb{S} \setminus \mathbb{V}}((\eta, \omega); A).$$

Since $\omega \in \Omega$ and $\mathbb{S} \in \mathcal{P}(\mathbb{T})$ were arbitrary, we conclude that $K_{\mathbb{V}}$ strongly preserves A by applying Eq. (4.4). $\square$

As a contrapositive of Lemma 4.3.4, if $\mathbb{V} \subseteq \mathbb{U}$ and $K_{\mathbb{V}}$ does not strongly preserve A , then $K_{\mathbb{U}}$ also does not strongly preserve A .

For the notions of causal effects and sources in Chapter 3, we developed extensions to the conditional and relative settings. For preservation, the conditional extension is

not meaningful, since conditioning on a delta measure simply yields the delta measure back. However, preservation can be generalised to relative preservation. That is to say, a causal kernel $K_{\mathbb{U}}$ may not preserve an event A by itself, but relative to another causal kernel $K_{\mathbb{W}}$, it may not have any further effect beyond what is encoded in $K_{\mathbb{W}}$. This is another way that relationships between causal kernels may be encoded, and is formalised in the following.

Definition 4.3.5 (Relative preservation). We say that the causal kernel $K_{\mathbb{U}}$ weakly preserves an event A *relative* to another causal kernel $K_{\mathbb{W}}$ if, for all $\eta, \omega \in \Omega$,

$$K_{\mathbb{U}\mathbb{W}}((\eta, \omega); A) = K_{\mathbb{W}}((\eta, \omega); A). \quad (4.5)$$

We say that $K_{\mathbb{U}}$ strongly preserves A relative to $K_{\mathbb{W}}$ if, for all $\mathbb{S} \in \mathcal{P}(\mathbb{T})$ and all $\eta, \omega \in \Omega$,

$$K_{\mathbb{S}\mathbb{U}\mathbb{W}}((\eta, \omega); A) = K_{(\mathbb{S}\mathbb{U}\mathbb{W}) \setminus (\mathbb{U} \setminus \mathbb{W})}((\eta, \omega); A). \quad (4.6)$$

For a sub- σ -algebra $\mathcal{F} \subseteq \mathcal{H}$, we say that $K_{\mathbb{U}}$ weakly/strongly preserves $\mathcal{F}$ relative to $K_{\mathbb{W}}$ if it weakly/strongly preserves every $A \in \mathcal{F}$ relative to $K_{\mathbb{W}}$.

If $\mathbb{W} = \emptyset$, then we trivially recover the (non-relative) preservation in Definition 4.3.1. We leave it as an exercise to prove results that are analogous to Lemmas 4.3.2 and 4.3.4 in the case of relative preservation, i.e. if $K_{\mathbb{U}}$ strongly preserves A relative to $K_{\mathbb{W}}$, then it weakly preserves A relative to $K_{\mathbb{W}}$, and if $K_{\mathbb{U}}$ strongly preserves A relative to $K_{\mathbb{W}}$, then for all $\mathbb{V} \subseteq \mathbb{U}$, $K_{\mathbb{V}}$ strongly preserves A relative to $K_{\mathbb{W}}$.

We end this section by making a couple of simple remarks that are specific to relative preservation. First, by letting $\mathbb{W} = \mathbb{U}$ in Definition 4.3.5, we always have, for all $\mathbb{S} \in \mathcal{P}(\mathbb{T})$ and all $\eta, \omega \in \Omega$,

$$K_{\mathbb{S}\mathbb{U}\mathbb{U}}((\eta, \omega); A) = K_{(\mathbb{S}\mathbb{U}\mathbb{U}) \setminus (\mathbb{U} \setminus \mathbb{U})}((\eta, \omega); A),$$

so $K_{\mathbb{U}}$ always strongly preserves any event A relative to itself.

Lastly, it is possible for $K_{\mathbb{U}}$ not to preserve A (weakly or strongly), yet for there to exist $\mathbb{W} \in \mathcal{P}(\mathbb{T})$ such that $K_{\mathbb{U}}$ preserves A (weakly or strongly) relative to $K_{\mathbb{W}}$. However, the converse is only true for strong preservation—it is possible for $K_{\mathbb{U}}$ to weakly preserve A , yet for there to exist $\mathbb{W} \in \mathcal{P}(\mathbb{T})$ such that $K_{\mathbb{U}}$ does not weakly preserve A relative to $K_{\mathbb{W}}$. This is again due to the fact that, a priori, there is no relationship between $K_{\mathbb{U}}$ and $K_{\mathbb{U}\mathbb{W}}$. On the other hand, if $K_{\mathbb{U}}$ strongly preserves A (in the sense of Definition 4.3.1), then Lemma 4.3.4 tells us that, for any $\mathbb{W} \in \mathcal{P}(\mathbb{T})$, the causal kernel $K_{\mathbb{U}\mathbb{W}}$ also strongly preserves A . Hence, for all $\mathbb{S} \in \mathcal{P}(\mathbb{T})$,

$$K_{\mathbb{S}\mathbb{U}\mathbb{W}}((\eta, \omega); A) = K_{(\mathbb{S}\mathbb{U}\mathbb{W}) \setminus (\mathbb{U} \setminus \mathbb{W})}((\eta, \omega); A),$$

by using $\mathbb{S}\mathbb{U}\mathbb{W}$ in the place of $\mathbb{S}$ in Eq. (4.4). But this is precisely Eq. (4.6), which means that $K_{\mathbb{U}}$ strongly preserves A relative to $K_{\mathbb{W}}$. Hence, (non-relative) strong preservation implies strong preservation relative to all other kernels $K_{\mathbb{W}}$.




Figure 4.6: Strong relative preservation (Definition 4.3.5). If fixing any $(\mathbb{S} \cup \mathbb{W})$ -coordinates to any value $\omega_{\mathbb{S} \cup \mathbb{W}}$ for any pre-intervention outcome $\eta \in \Omega$ is the same as fixing the $((\mathbb{S} \cup \mathbb{W}) \setminus (\mathbb{U} \setminus \mathbb{W}))$ -coordinates to the same (marginal) values $\omega_{(\mathbb{S} \cup \mathbb{W}) \setminus (\mathbb{U} \setminus \mathbb{W})}$ (i.e. $K_{\mathbb{U}}((\eta, \omega); A) = K_{(\mathbb{S} \cup \mathbb{W}) \setminus (\mathbb{U} \setminus \mathbb{W})}((\eta, \omega); A)$ for all $\mathbb{S} \in \mathcal{P}(\mathbb{T})$ and all $\eta, \omega \in \Omega$), then $K_{\mathbb{U}}$ strongly preserves A .

4.4 Targeted causal kernels

So far, the discussion has focused on causal kernels $K_{\mathbb{U}}$ for fixed coordinates $\mathbb{U} \in \mathcal{P}(\mathbb{T})$. However, it is sometimes of interest to intervene on different coordinates altogether depending on the pre-intervention information. This could be of interest, for example, if we have separate coordinates representing the number of physiotherapy sessions and the dose of a medicine, and depending on the age of the patient, we either assign them to physiotherapy or medicine, or both. In the study of a stochastic process, one may want to intervene once the process reaches a threshold, i.e. we may want to intervene at a random stopping time, rather than at a fixed time point.

To this end, we define a causal kernel that is not fixed at a particular set of coordinates $\mathbb{U}$, but one that varies according to a *targeting function*. The purpose of the careful measurability set-up at the beginning of Section 4.1, with the σ -algebra $\mathcal{T}$ on the power set $\mathcal{P}(\mathbb{T})$ (Eq. (4.1)) and the σ -algebra $\mathcal{U}$ on $\mathcal{P}(\mathbb{T}) \times \Omega \times \Omega$ (Eq. (4.2)), was precisely to enable such targeting. We start by defining the *targeted* causal kernel as follows.

Definition 4.4.1 (Targeted causal kernel). Suppose that $U : \Omega \rightarrow \mathcal{P}(\mathbb{T})$ is a $\mathcal{H}/\mathcal{T}$ -measurable map, called the *targeting function*. We define the *targeted causal kernel* as the map $K_U : \Omega \times \Omega \times \mathcal{H} \rightarrow [0, 1]$ given by

$$K_U((\eta, \omega); A) = K_{U(\eta)}((\eta, \omega); A).$$

Semantically, for each pre-intervention outcome η , the targeting function U first iden-




Figure 4.7: The targeted causal kernel $K_U((\eta, \omega); \cdot)$ (Definition 4.4.1) is the measure when, for the pre-intervention outcome η , we fix the $U(\eta)$ -coordinates to the values $\omega_{U(\eta)}$. The coordinates which get fixed during an intervention depend on the pre-intervention outcome η via the targeting function U .

tifies the coordinates $U(\eta)$ to intervene on. The targeted causal kernel $K_U((\eta, \omega); A)$ then tells us the measure we obtain on an event A when fixing those coordinates to $\omega_{U(\eta)}$. See Fig. 4.7 for an illustration.

The pre-intervention information on which the targeting operation is based is formalised as follows.

Definition 4.4.2 (Targeting scope). We define the *scope* of the targeting function $U : \Omega \rightarrow \mathcal{P}(\mathbb{T})$, denoted by τ_U , as the sub- σ -algebra of $\mathcal{H}$ generated by the function U . Formally,

$$\tau_U = \sigma(U) = U^{-1}(\mathcal{T}) \subseteq \mathcal{H}.$$

If U is a constant function and takes every η to the same set of coordinates $\mathbb{U} \in \mathcal{P}(\mathbb{T})$, then the targeting scope is trivial: $\tau_U = \{\emptyset, \Omega\}$.

In the case of fixed coordinates $\mathbb{U}$, the causal kernel $K_{\mathbb{U}}$ encodes information about interventions on $\mathcal{H}_{\mathbb{U}}$, the sub- σ -algebra of $\mathcal{H}$ representing those coordinates. Unfortunately, for a targeted kernel K_U , we do not have a fixed corresponding sub- σ -algebra of $\mathcal{H}$. We can, however, define a σ -algebra $\mathcal{H}_U$ on the augmented space $\Omega \times \Omega$, with the first Ω for the pre-intervention outcomes η and the second Ω for the outcomes ω with which we intervene, as follows:

$$\mathcal{H}_U = \sigma(\{(\eta, \omega) : t \in U(\eta), \omega \in H\} : t \in \mathbb{T}, H \in \mathcal{H}_t). \quad (4.7)$$

We will call it the “targeted coordinate σ -algebra” when it must be called by a name, but most often, we will simply refer to it by its notation $\mathcal{H}_U$. Semantically, $\mathcal{H}_U$ encodes the pre-intervention information on which the targeting is based (i.e. the targeting scope—Definition 4.4.2) in the first component η , and in the second component, what outcome ω was used for the intervention on the targeted coordinates $U(\eta)$. Each generating set

$$\{(\eta, \omega) : t \in U(\eta), \omega \in H\}$$

for $t \in \mathbb{T}$ and $H \in \mathcal{H}_t$ is a measurable set in $\mathcal{H} \otimes \mathcal{H}$, since the σ -algebra $\mathcal{T}$ on $\mathcal{P}(\mathbb{T})$ is defined so that each $\{\mathbb{S} : t \in \mathbb{S}\} \subseteq \mathcal{P}(\mathbb{T})$ is in $\mathcal{T}$, and by the $\mathcal{H}/\mathcal{T}$ -measurability of U , the set $\{\eta \in \Omega : t \in U(\eta)\} = U^{-1}(\{\mathbb{S} : t \in \mathbb{S}\})$ is measurable.

If $U(\eta) = \mathbb{U}$ is constant (i.e. if the scope τ_U is trivial), then $\mathcal{H}_U = \{\emptyset, \Omega\} \otimes \mathcal{H}_{\mathbb{U}}$, i.e. it is the pullback of $\mathcal{H}_{\mathbb{U}}$ under the projection map $(\eta, \omega) \mapsto \omega : \Omega \times \Omega \rightarrow \Omega$. This is because the generating sets reduce to

$$\{(\eta, \omega) : \omega \in H\}$$

for $H \in \mathcal{H}_t$ if $t \in \mathbb{U}$, and nothing if $t \notin \mathbb{U}$. The sets $\{\eta : t \in U(\eta)\}$ are the entire Ω if $t \in \mathbb{U}$ and the empty set $\emptyset$ otherwise.

The targeted coordinate σ -algebra $\mathcal{H}_U$ (Eq. (4.7)) can be characterised using *sections*, which we recall from Eq. (2.1). For a set $D \in \mathcal{H} \otimes \mathcal{H}$ and any pre-intervention outcome η , let us denote the η -section of D as

$$D_\eta := \{\omega \in \Omega : (\eta, \omega) \in D\}. \quad (4.8)$$

Then we have the following characterisation of $\mathcal{H}_U$.

Lemma 4.4.3. *For each $\eta \in \Omega$, we have $\{D_\eta : D \in \mathcal{H}_U\} = \mathcal{H}_{U(\eta)}$.*

Proof. Fix $\eta \in \Omega$. We first show that $D_\eta \in \mathcal{H}_{U(\eta)}$ for every $D \in \mathcal{H}_U$. To this end, let us define

$$\mathcal{D}_\eta := \{D \subseteq \Omega \times \Omega : D_\eta \in \mathcal{H}_{U(\eta)}\}.$$

Since sectioning commutes with complements and countable unions, namely

$$(\Omega \times \Omega \setminus D)_\eta = \Omega \setminus D_\eta \quad \text{and} \quad (\cup_{n \in \mathbb{N}} D_n)_\eta = \cup_{n \in \mathbb{N}} (D_n)_\eta,$$

the class $\mathcal{D}_\eta$ is a σ -algebra on $\Omega \times \Omega$.

Now take any generator

$$G_{t,H} := \{(\eta', \omega) : t \in U(\eta'), \omega \in H\}, \quad t \in \mathbb{T}, H \in \mathcal{H}_t$$

of $\mathcal{H}_U$ (cf. Eq. (4.7)). Its η -section is

$$\begin{aligned} (G_{t,H})_\eta &= \{\omega \in \Omega : (\eta, \omega) \in G_{t,H}\} \\ &= \{\omega \in \Omega : t \in U(\eta), \omega \in H\} \\ &= \begin{cases} H & \text{if } t \in U(\eta), \\ \emptyset & \text{if } t \notin U(\eta). \end{cases} \end{aligned}$$

If $t \in U(\eta)$, then $H \in \mathcal{H}_t \subseteq \mathcal{H}_{U(\eta)}$; if $t \notin U(\eta)$, then $\emptyset \in \mathcal{H}_{U(\eta)}$. Therefore $(G_{t,H})_\eta \in \mathcal{H}_{U(\eta)}$, so each generator $G_{t,H}$ belongs to $\mathcal{D}_\eta$. Since $\mathcal{D}_\eta$ is a σ -algebra containing all generators of $\mathcal{H}_U$, we obtain $\mathcal{H}_U \subseteq \mathcal{D}_\eta$, i.e.

$$\{D_\eta : D \in \mathcal{H}_U\} \subseteq \mathcal{H}_{U(\eta)}.$$

For the reverse inclusion, note that $\mathcal{H}_{U(\eta)}$ is generated by the sets $H \in \mathcal{H}_t$ with $t \in U(\eta)$. If $t \in U(\eta)$ and $H \in \mathcal{H}_t$, then for the generator $G_{t,H} \in \mathcal{H}_U$ we have

$$(G_{t,H})_\eta = H.$$

Hence every generator of $\mathcal{H}_{U(\eta)}$ belongs to the family $\{D_\eta : D \in \mathcal{H}_U\}$. Finally, the family $\{D_\eta : D \in \mathcal{H}_U\}$ is itself a σ -algebra on Ω , again because sectioning commutes with complements and countable unions. Since it contains all generators of $\mathcal{H}_{U(\eta)}$, it follows that

$$\mathcal{H}_{U(\eta)} \subseteq \{D_\eta : D \in \mathcal{H}_U\}.$$

Combining the two inclusions yields

$$\{D_\eta : D \in \mathcal{H}_U\} = \mathcal{H}_{U(\eta)},$$

as required. $\square$

In the next result, we show that the targeted causal kernel K_U is a transition probability kernel with respect to the σ -algebra

$$\tilde{\mathcal{H}}_U := (\mathcal{H} \otimes \{\emptyset, \Omega\}) \vee \mathcal{H}_U. \quad (4.9)$$

The σ -algebra $\tilde{\mathcal{H}}_U$ contains full information on the η -component (rather than just the targeting scope, as $\mathcal{H}_U$ does), and only the information about the coordinates chosen by the targeting function U for that particular η on the ω -component. Since causation may be heterogeneous within each set of targeted coordinates in $\text{Im}(U)$, the σ -algebra with respect to which K_U is measurable has to be enlarged to take this information into account.

Recall that $\mathcal{U}$ (Eq. (4.2)) is a σ -algebra on $\mathcal{P}(\mathbb{T}) \times \Omega \times \Omega$ that contains full information on the $(\mathbb{S}, \eta)$ components, but only the information included in the $\mathcal{H}_\mathbb{S}$ coordinates for the ω -component. The σ -algebra $\tilde{\mathcal{H}}_U$ can be seen as a targeted version of $\mathcal{U}$, by substituting in $U(\eta)$ for $\mathbb{S}$. In other words, $\mathcal{U}$ describes what information is available before the coordinates are chosen, while $\tilde{\mathcal{H}}_U$ describes the corresponding information after the choice $\mathbb{S} = U(\eta)$ has been made through the targeting function U .

Theorem 4.4.4. *K_U is a transition probability kernel from $(\Omega \times \Omega, \tilde{\mathcal{H}}_U)$ into $(\Omega, \mathcal{H})$.*

Proof. For a fixed pair $(\eta, \omega) \in \Omega \times \Omega$, the map $A \mapsto K_{U(\eta)}((\eta, \omega); A) : \mathcal{H} \rightarrow [0, 1]$ is clearly a probability measure, since the definition of the causal mechanism ensures that for each fixed $\mathbb{S} \in \mathcal{P}(\mathbb{T})$ and $\eta, \omega \in \Omega$, the map $A \mapsto K_\mathbb{S}((\eta, \omega); A)$ is a probability measure. It remains to show that, for each $A \in \mathcal{H}$, the map $(\eta, \omega) \mapsto K_{U(\eta)}((\eta, \omega); A)$ is $\tilde{\mathcal{H}}_U$ -measurable.

The definition of the causal mechanism $\mathbb{K}$ (Definition 4.1.1) ensures that, for each event $A \in \mathcal{H}$, the map $(\mathbb{S}, \eta, \omega) \mapsto K_{\mathbb{S}}((\eta, \omega); A)$ is $\mathcal{U}$ -measurable. The map $(\eta, \omega) \mapsto K_{U(\eta)}((\eta, \omega); A)$ can be seen as a composition of the map $\iota_U : \Omega \times \Omega \rightarrow \mathcal{P}(\mathbb{T}) \times \Omega \times \Omega$ defined by

$$\iota_U(\eta, \omega) = (U(\eta), \eta, \omega) \quad (4.10)$$

and the $\mathcal{U}$ -measurable map $(\mathbb{S}, \eta, \omega) \mapsto K_{\mathbb{S}}((\eta, \omega); A)$. It now suffices to show that ι_U is $\tilde{\mathcal{H}}_U/\mathcal{U}$ -measurable.

The σ -algebra $\mathcal{U}$ is generated by sets of the form $D \times E \times \Omega$, with $D \in \mathcal{T}$ and $E \in \mathcal{H}$, together with sets of the form

$$\{(\mathbb{S}, \eta, \omega) : t \in \mathbb{S}, \omega \in H\}, t \in \mathbb{T}, H \in \mathcal{H}_t.$$

For the first type,

$$\iota_U^{-1}(D \times E \times \Omega) = (U^{-1}(D) \cap E) \times \Omega \in \mathcal{H} \otimes \{\emptyset, \Omega\} \subseteq \tilde{\mathcal{H}}_U,$$

since U is $\mathcal{H}/\mathcal{T}$ -measurable. For the second type,

$$\iota_U^{-1}(\{(\mathbb{S}, \eta, \omega) : t \in \mathbb{S}, \omega \in H\}) = \{(\eta, \omega) : t \in U(\eta), \omega \in H\} \in \mathcal{H}_U \subseteq \tilde{\mathcal{H}}_U.$$

Hence ι_U is $\tilde{\mathcal{H}}_U/\mathcal{U}$ -measurable, and as a result, the map $(\eta, \omega) \mapsto K_{U(\eta)}((\eta, \omega); A)$ is $\tilde{\mathcal{H}}_U$ -measurable. Therefore K_U is a transition probability kernel from $(\Omega \times \Omega, \tilde{\mathcal{H}}_U)$ into $(\Omega, \mathcal{H})$. $\square$

In fact, we have $\tilde{\mathcal{H}}_U = \iota_U^{-1}(\mathcal{U})$. Indeed, the generators of $\tilde{\mathcal{H}}_U$ are $E \times \Omega$ for $E \in \mathcal{H}$, and $\{(\eta, \omega) : t \in U(\eta), \omega \in H\}$ for $t \in \mathbb{T}$ and $H \in \mathcal{H}_t$. For the first type, we can take $\mathcal{P}(\mathbb{T}) \times E \times \Omega \in \mathcal{U}$ and see that

$$\iota_U^{-1}(\mathcal{P}(\mathbb{T}) \times E \times \Omega) = E \times \Omega,$$

and for the second type, we can take $\{(\mathbb{S}, \eta, \omega) : t \in \mathbb{S}, \omega \in H\} \in \mathcal{U}$ and see (as we did in the proof above) that

$$\iota_U^{-1}(\{(\mathbb{S}, \eta, \omega) : t \in \mathbb{S}, \omega \in H\}) = \{(\eta, \omega) : t \in U(\eta), \omega \in H\}.$$

Unlike the causal kernels $K_{\mathbb{U}}$ at fixed coordinates $\mathbb{U} \in \mathcal{P}(\mathbb{T})$, the targeted causal kernels K_U are not part of the causal mechanism $\mathbb{K}$ —they are only defined in tandem with a targeting function U . However, most of the definitions from the previous sections carry over seamlessly to the targeted case. Perhaps the most fundamental is a targeted version of the interventional determinism axiom for K_U : for all $\eta, \omega \in \Omega$ and all $B \in \mathcal{H}_{U(\eta)}$,

$$K_U((\eta, \omega); B) = \mathbf{1}_B(\omega).$$

The proof of this property is immediate: we can apply the interventional determinism axiom for each $U(\eta)$.

The notion of heterogeneity requires a little bit more care when carrying over from the fixed-coordinate case than the definitions above—careful adjustments must be made

using $\mathcal{H}_U$ as an analogue of $\mathcal{H}_\mathbb{U}$. For any sub- σ -algebra $\mathcal{J} \subseteq \mathcal{H}$, the pullback of $\mathcal{J}$ under the projection $(\eta, \omega) \mapsto \eta$ is $\mathcal{J} \otimes \{\emptyset, \Omega\}$. In other words,

$$\mathcal{J} \otimes \{\emptyset, \Omega\} = \{B \times \Omega : B \in \mathcal{J}\} \subseteq \mathcal{H} \otimes \mathcal{H}.$$

Then we can augment $\mathcal{H}_U$ (Eq. (4.7)) with the information in $\mathcal{J}$ to obtain a σ -algebra $\tilde{\mathcal{H}}_U^\mathcal{J}$ on $\Omega \times \Omega$ as follows:

$$\tilde{\mathcal{H}}_U^\mathcal{J} = (\mathcal{J} \otimes \{\emptyset, \Omega\}) \vee \mathcal{H}_U \subseteq \mathcal{H} \otimes \mathcal{H}. \quad (4.11)$$

This is reminiscent of how we defined $\tilde{\mathcal{H}}_U$ in Eq. (4.9), but by augmenting $\mathcal{H}_U$ with general sub- σ -algebras $\mathcal{J} \subseteq \mathcal{H}$ rather than the entire $\mathcal{H}$, to allow for heterogeneity within each targeted coordinate $\mathbb{U} \in \text{Im}(U)$. Let us now define the heterogeneity family of K_U , analogously to Definition 4.2.1.

Definition 4.4.5 (Heterogeneity family—targeted kernels). For each event $A \in \mathcal{H}$, a σ -algebra $\mathcal{J} \subseteq \mathcal{H}$ is a *heterogeneity σ -algebra* of the targeted causal kernel K_U if the map $(\eta, \omega) \mapsto K_U((\eta, \omega); A)$ is $\tilde{\mathcal{H}}_U^\mathcal{J}$ -measurable. We then define the heterogeneity family of K_U on A as the collection of all of its heterogeneity σ -algebras on A , and denote it by $\Sigma_U(A)$.

For a sub- σ -algebra $\mathcal{F} \subseteq \mathcal{H}$, a σ -algebra $\mathcal{J} \subseteq \mathcal{H}$ is a heterogeneity σ -algebra of K_U on $\mathcal{F}$ if $\mathcal{J} \in \Sigma_U(A)$ for every $A \in \mathcal{F}$. We define the heterogeneity family of K_U on $\mathcal{F}$ as

$$\Sigma_U(\mathcal{F}) = \bigcap_{A \in \mathcal{F}} \Sigma_U(A);$$

in other words, it is the collection of all of its heterogeneity σ -algebras on $\mathcal{F}$.

If $\mathcal{J}$ is a heterogeneity σ -algebra of K_U on $\mathcal{H}$, we simply say that it is a heterogeneity σ -algebra of K_U . We also simply write $\Sigma_U = \Sigma_U(\mathcal{H})$.

Note that, because the map $(\eta, \omega) \mapsto K_U((\eta, \omega); A)$ is $\tilde{\mathcal{H}}_U$ -measurable by Theorem 4.4.4, $\mathcal{H}$ is always a heterogeneity σ -algebra, hence a heterogeneity family of a targeted causal kernel is never empty. In other words, for all events $A \in \mathcal{H}$, we have $\mathcal{H} \in \Sigma_U(A)$; as a consequence, we also have $\mathcal{H} \in \Sigma_U(\mathcal{F})$ for any $\mathcal{F} \subseteq \mathcal{H}$. Just like Definition 4.2.1, we have that a heterogeneity family is upward closed: for an event $A \in \mathcal{H}$, if $\mathcal{J} \in \Sigma_U(A)$, then for all $\tilde{\mathcal{J}} \supset \mathcal{J}$, we have $\tilde{\mathcal{J}} \in \Sigma_U(A)$.

As in the fixed-coordinate case (cf. Remark 4.2.2), it is necessary to work with families of heterogeneity σ -algebras, rather than looking for a single minimum σ -algebra containing the heterogeneity information. The issue is again that $\bigcap_{\mathcal{J} \in \Sigma_U(A)} \tilde{\mathcal{H}}_U^\mathcal{J}$ may not equal $\tilde{\mathcal{H}}_U^{\bigcap_{\mathcal{J} \in \Sigma_U(A)} \mathcal{J}}$, unless additional regularity assumptions are imposed.

If $U(\eta) = \mathbb{U}$ is a constant function, then we recover Definition 4.2.1, by the above argument that $\mathcal{H}_U$ is the pullback of $\mathcal{H}_\mathbb{U}$ under the projection $(\eta, \omega) \mapsto \omega$. Specifically, we have

$$\tilde{\mathcal{H}}_U^\mathcal{J} = (\mathcal{J} \otimes \{\emptyset, \Omega\}) \vee (\{\emptyset, \Omega\} \otimes \mathcal{H}_\mathbb{U}) = \mathcal{J} \otimes \mathcal{H}_\mathbb{U},$$

so $\Sigma_U(A)$, $\Sigma_U(\mathcal{F})$, Σ_U in this case reduce immediately to $\Sigma_\mathbb{U}(A)$, $\Sigma_\mathbb{U}(\mathcal{F})$, $\Sigma_\mathbb{U}$ given in Definition 4.2.1.

The notion of *homogeneous* targeted kernels is defined analogously to the non-targeted case.

Definition 4.4.6 (Targeted homogeneous kernels). A targeted causal kernel K_U is *homogeneous* on A if $\{\emptyset, \Omega\} \in \Sigma_U(A)$, and heterogeneous otherwise. It is homogeneous on $\mathcal{F}$ if $\{\emptyset, \Omega\} \in \Sigma_U(\mathcal{F})$, and heterogeneous otherwise. If $\{\emptyset, \Omega\} \in \Sigma_U$, then we simply say that K_U is homogeneous.

Unlike non-targeted homogeneous kernels, which have no dependence on the pre-intervention argument η , targeted homogeneous kernels are still dependent on the pre-intervention argument η through the targeting operation. Homogeneity only means that, once the targeting coordinates $U(\eta)$ are chosen, no further pre-intervention information is needed.

The scope τ_U is generated by the sets $\{\eta : t \in U(\eta)\}$, $t \in \mathbb{T}$, and by the definition of $\mathcal{H}_U$ (Eq. (4.7)), we have that

$$\{\eta : t \in U(\eta)\} \times \Omega = \{(\eta, \omega) : t \in U(\eta), \omega \in \Omega\} \in \mathcal{H}_U,$$

i.e. $\tau_U \otimes \{\emptyset, \Omega\} \subseteq \mathcal{H}_U$. Hence, if $\mathcal{J}$ is contained in the targeting scope, $\mathcal{J} \subseteq \tau_U$, then we have $\tilde{\mathcal{H}}_U^{\mathcal{J}} = \mathcal{H}_U$. Accordingly, homogeneity of a targeted causal kernel K_U on A is equivalent to $\tau_U \in \Sigma_U(A)$, since this is sufficient to imply that no further η -variability exists beyond what is encoded in $\mathcal{H}_U$ for targeting. In other words, the map $(\eta, \omega) \mapsto K_U((\eta, \omega); A)$ is $\mathcal{H}_U$ -measurable. Likewise, the homogeneity of K_U on a σ -algebra $\mathcal{F}$ is equivalent to $\tau_U \in \Sigma_U(\mathcal{F})$, and the homogeneity of K_U to $\tau_U \in \Sigma_U$. Hence, the homogeneity here is within each targeted coordinate $\mathbb{U} \in \text{Im}(U)$.

In the same vein, we define the homogenisation $\overline{K}_U$ of K_U , for $A \in \mathcal{H}$, as

$$\overline{K}_U((\eta, \omega); A) = \overline{K}_{U(\eta)}(\omega; A) = \int \mathbb{P}(d\eta') K_{U(\eta)}((\eta', \omega); A). \quad (4.12)$$

Unlike the homogenisation $\overline{K}_{\mathbb{U}}$ of causal kernel $K_{\mathbb{U}}$ on fixed coordinates $\mathbb{U}$ (Definition 4.2.3), the homogenised targeted kernel $\overline{K}_U$ retains its pre-intervention η -argument. However, the η argument is again only used for targeting; within each $U(\eta)$, we integrate out the heterogeneity argument η' in $K_{U(\eta)}$ with respect to the observational measure $\mathbb{P}$. If the targeting function is constant, i.e. $U(\eta) = \mathbb{U}$ for some $\mathbb{U} \in \mathcal{P}(\mathbb{T})$, then we immediately recover $\overline{K}_U((\eta, \omega); A) = \overline{K}_{\mathbb{U}}(\omega; A)$ for all $\eta, \omega \in \Omega$ and all $A \in \mathcal{H}$.

Next up is the targeted analogue of preservation (Definition 4.3.1). To iron out the measurability issues arising from using multiple targeting functions, we start with the following simple lemma.

Lemma 4.4.7. *If $U, W : \Omega \rightarrow \mathcal{P}(\mathbb{T})$ are $\mathcal{H}/\mathcal{T}$ -measurable, then so are $U \cup W$, $U \cap W$, and $U \setminus W$, given by $(U \cup W)(\eta) = U(\eta) \cup W(\eta)$, $(U \cap W)(\eta) = U(\eta) \cap W(\eta)$ and $(U \setminus W)(\eta) = U(\eta) \setminus W(\eta)$ for each $\eta \in \Omega$.*

Proof. Recall from Eq. (4.1) that $\mathcal{T}$ is generated by the sets $\{\mathbb{S} \in \mathcal{P}(\mathbb{T}) : t \in \mathbb{S}\}$ with $t \in \mathbb{T}$. Hence, to prove that a map $V : \Omega \rightarrow \mathcal{P}(\mathbb{T})$ is $\mathcal{H}/\mathcal{T}$ -measurable, it suffices to show that $V^{-1}(\{\mathbb{S} \in \mathcal{P}(\mathbb{T}) : t \in \mathbb{S}\}) \in \mathcal{H}$ for every $t \in \mathbb{T}$.

Now fix $t \in \mathbb{T}$. For the union map, we have

$$\begin{aligned}
 (U \cup W)^{-1}(\{\mathbb{S} \in \mathcal{P}(\mathbb{T}) : t \in \mathbb{S}\}) &= \{\eta \in \Omega : t \in (U \cup W)(\eta)\} \\
 &= \{\eta \in \Omega : t \in U(\eta) \cup W(\eta)\} \\
 &= \{\eta \in \Omega : t \in U(\eta)\} \cup \{\eta \in \Omega : t \in W(\eta)\} \\
 &= U^{-1}(\{\mathbb{S} \in \mathcal{P}(\mathbb{T}) : t \in \mathbb{S}\}) \cup W^{-1}(\{\mathbb{S} \in \mathcal{P}(\mathbb{T}) : t \in \mathbb{S}\}) \\
 &\in \mathcal{H},
 \end{aligned}$$

since U and W are $\mathcal{H}/\mathcal{T}$ -measurable.

Similarly, for the intersection map,

$$\begin{aligned}
 (U \cap W)^{-1}(\{\mathbb{S} \in \mathcal{P}(\mathbb{T}) : t \in \mathbb{S}\}) &= \{\eta \in \Omega : t \in (U \cap W)(\eta)\} \\
 &= \{\eta \in \Omega : t \in U(\eta) \cap W(\eta)\} \\
 &= \{\eta \in \Omega : t \in U(\eta)\} \cap \{\eta \in \Omega : t \in W(\eta)\} \\
 &= U^{-1}(\{\mathbb{S} \in \mathcal{P}(\mathbb{T}) : t \in \mathbb{S}\}) \cap W^{-1}(\{\mathbb{S} \in \mathcal{P}(\mathbb{T}) : t \in \mathbb{S}\}) \\
 &\in \mathcal{H}.
 \end{aligned}$$

Finally, for the set difference,

$$\begin{aligned}
 (U \setminus W)^{-1}(\{\mathbb{S} \in \mathcal{P}(\mathbb{T}) : t \in \mathbb{S}\}) &= \{\eta \in \Omega : t \in (U \setminus W)(\eta)\} \\
 &= \{\eta \in \Omega : t \in U(\eta) \setminus W(\eta)\} \\
 &= \{\eta \in \Omega : t \in U(\eta)\} \cap \{\eta \in \Omega : t \notin W(\eta)\} \\
 &= U^{-1}(\{\mathbb{S} \in \mathcal{P}(\mathbb{T}) : t \in \mathbb{S}\}) \cap (\Omega \setminus W^{-1}(\{\mathbb{S} \in \mathcal{P}(\mathbb{T}) : t \in \mathbb{S}\})) \\
 &\in \mathcal{H}.
 \end{aligned}$$

Thus $U \cup W$, $U \cap W$, and $U \setminus W$ are all $\mathcal{H}/\mathcal{T}$ -measurable. $\square$

Lemma 4.4.7 implies that, for fixed $\mathbb{S} \in \mathcal{P}(\mathbb{T})$, the maps $\mathbb{S} \cup U$ and $\mathbb{S} \setminus U$ are also $\mathcal{H}/\mathcal{T}$ -measurable, by choosing constant targeting functions. Now we present the definition of preservation for targeted kernels.

Definition 4.4.8 (Preservation—targeted kernels). The targeted causal kernel K_U weakly preserves an event A if, for all $\eta, \omega \in \Omega$, we have

$$K_U((\eta, \omega); A) = \mathbf{1}_A(\eta) = \delta_\eta(A).$$

We say that K_U strongly preserves an event A if, for all $\mathbb{S} \in \mathcal{P}(\mathbb{T})$ and all $\eta, \omega \in \Omega$, we have

$$K_{\mathbb{S}}((\eta, \omega); A) = K_{(\mathbb{S} \setminus U)}((\eta, \omega); A).$$

We say that K_U weakly/strongly preserves a sub- σ -algebra $\mathcal{F} \subseteq \mathcal{H}$ if it weakly/strongly preserves all $A \in \mathcal{F}$.

Note that, in strong preservation, we quantified over all possible coordinates $\mathbb{S} \in \mathcal{P}(\mathbb{T})$ as we did in Definition 4.3.1, rather than over all possible targeting functions $S : \Omega \rightarrow \mathcal{P}(\mathbb{T})$, which appears stronger and seems more natural for targeted kernels. However, it turns out that they are equivalent, as the next lemma shows.

Lemma 4.4.9. *A targeted causal kernel K_U strongly preserves an event A if and only if, for all $\mathcal{H}/\mathcal{T}$ -measurable maps $S : \Omega \rightarrow \mathcal{P}(\mathbb{T})$ and all $\eta, \omega \in \Omega$, we have*

$$K_S((\eta, \omega); A) = K_{S \setminus U}((\eta, \omega); A).$$

Proof. Sufficiency is obvious, by considering constant maps $S(\eta) = \mathbb{S}$ for each $\mathbb{S}$. For necessity, suppose that K_U strongly preserves A . Then for any $\mathcal{H}/\mathcal{T}$ -measurable targeting function $S : \Omega \rightarrow \mathcal{P}(\mathbb{T})$ and any $\eta, \omega \in \Omega$, we have

$$\begin{aligned} K_S((\eta, \omega); A) &= K_{S(\eta)}((\eta, \omega); A) \\ &= K_{S(\eta) \setminus U(\eta)}((\eta, \omega); A) \quad \text{since } K_U \text{ strongly preserves } A \\ &= K_{S \setminus U}((\eta, \omega); A), \end{aligned}$$

as required. $\square$

For relative preservation of targeted kernels given in the next definition (cf. Definition 4.3.5), not only is the main kernel K_U targeted, but the auxiliary kernel K_W is also targeted.

Definition 4.4.10 (Relative preservation—targeted kernels). We say that the targeted causal kernel K_U weakly preserves an event A relative to another targeted causal kernel K_W if, for all $\eta, \omega \in \Omega$,

$$K_{U \cup W}((\eta, \omega); A) = K_W((\eta, \omega); A),$$

where the map $(U \cup W) : \Omega \rightarrow \mathcal{P}(\mathbb{T})$ is a map defined by $(U \cup W)(\eta) = U(\eta) \cup W(\eta)$. We say that K_U strongly preserves A relative to K_W if, for all $\mathbb{S} \in \mathcal{P}(\mathbb{T})$ and all $\eta, \omega \in \Omega$,

$$K_{\mathbb{S} \cup W}((\eta, \omega); A) = K_{(\mathbb{S} \cup W) \setminus (U \setminus W)}((\eta, \omega); A).$$

For a sub- σ -algebra $\mathcal{F} \subseteq \mathcal{H}$, we say that K_U weakly/strongly preserves $\mathcal{F}$ relative to K_W if it weakly/strongly preserves every $A \in \mathcal{F}$ relative to K_W .

Bibliography

- Bruce Alberts, Rebecca Heald, Julian Lewis, David Morgan, Martin Raff, Keith Roberts, and Peter Walter. *Molecular Biology of the Cell*. W. W. Norton & Co., 1983.
- John Alcock. *Animal Behavior: An Evolutionary Approach*. Sinauer associates, 2009.
- Roy M Anderson and Robert M May. *Infectious Diseases of Humans: Dynamics and Control*. Oxford university press, 1991.
- Joshua D Angrist and Jörn-Steffen Pischke. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton university press, 2009.
- Peter William Atkins, Julio De Paula, and James Keeler. *Atkins' Physical Chemistry*. Oxford university press, 2023.
- Fausto Barbero. On the Logic of Interventionist Counterfactuals under Indeterministic Causal Laws. In *International Symposium on Foundations of Information and Knowledge Systems*, pages 203–221. Springer, 2024.
- Fausto Barbero and Gabriel Sandu. Team Semantics for Interventionist Counterfactuals: Observations vs. Interventions. *Journal of Philosophical Logic*, 50(3):471–521, 2021.
- Fausto Barbero and Jonni Virtema. Strongly Complete Axiomatization for a Logic with Probabilistic Interventionist Counterfactuals. In *European Conference on Logics in Artificial Intelligence*, pages 649–664. Springer, 2023.
- Fausto Barbero and Fan Yang. Characterizing Counterfactuals and Dependencies over (Generalized) Causal Teams. *Notre Dame Journal of Formal Logic*, 63(3):301–341, 2022.
- Fausto Barbero, Katrin Schulz, Fernando R Velázquez-Quesada, and Kaibo Xie. Observing Interventions: a Logic for Thinking about Experiments. *Journal of Logic and Computation*, 33(6):1152–1185, 2023.
- Sander Beckers. Nondeterministic Causal Models. In *Causal Learning and Reasoning*, pages 1532–1554. PMLR, 2025.
- Sander Beckers, Joseph Halpern, and Christopher Hitchcock. Causal Models with Constraints. In *Conference on causal learning and reasoning*, pages 866–879. PMLR, 2023.
- Michael Begon and Colin R Townsend. *Ecology: From Individuals to Ecosystems*. John Wiley & Sons, 2020.

- Patrick Billingsley. *Probability and Measure*. John Wiley & Sons, 1976.
- Philip Boeken and Joris M Mooij. Dynamic Structural Causal Models. *arXiv preprint arXiv:2406.01161*, 2024.
- Blai Bonet. A Calculus for Causal Relevance. In *Proceedings of the Seventeenth conference on Uncertainty in artificial intelligence*, pages 40–47, 2001.
- Stephan Bongers, Tineke Blom, and Joris M Mooij. Causal Modeling of Dynamical Systems. *arXiv preprint arXiv:1803.08784*, 2018.
- Stephan Bongers, Patrick Forré, Jonas Peters, and Joris M Mooij. Foundations of Structural Causal Models with Cycles and Latent Variables. *The Annals of Statistics*, 49(5):2885–2915, 2021.
- Émile Borel. Les Probabilités Dénombrables et leurs Applications Arithmétiques. *Rendiconti del Circolo Matematico di Palermo (1884-1940)*, 27(1):247–271, 1909.
- Allan M Brandt. Inventing Conflicts of Interest: a History of Tobacco Industry Tactics. *American journal of public health*, 102(1):63–71, 2012.
- Robert J Brulle. The Climate Lobby: A Sectoral Analysis of Lobbying Spending on Climate Change in the USA, 2000 to 2016. *Climatic change*, 149(3):289–303, 2018.
- Simon Buchholz, Junhyung Park, and Bernhard Schölkopf. Products, Abstractions and Inclusions of Causal Spaces. In *The 40th Conference on Uncertainty in Artificial Intelligence*, 2024.
- Daphna Buchsbaum, Sophie Bridgers, Deena Skolnick Weisberg, and Alison Gopnik. The Power of Possibility: Causal Learning, Counterfactual Reasoning, and Pretend Play. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 367(1599):2202–2212, 2012.
- Ruth MJ Byrne. Counterfactual Thought. *Annual review of psychology*, 67(1):135–157, 2016.
- Ruth MJ Byrne and Alice McEleney. Counterfactual Thinking about Actions and Failures to Act. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26(5):1318, 2000.
- Erhan Çinlar. *Probability and Stochastics*. Springer, 2011.
- John Collins, Ned Hall, and Laurie Ann Paul. *Causation and Counterfactuals*. MIT Press, 2004.
- A Philip Dawid. Who Needs Counterfactuals? In *Causal Models and Intelligent Data Management*, pages 33–50. Springer, 1999.
- A Philip Dawid. Causal Inference Without Counterfactuals. *Journal of the American statistical Association*, 95(450):407–424, 2000.

- A Philip Dawid. Counterfactuals, Hypotheticals and Potential Responses: A Philosophical Examination of Statistical Causality. *Research Report No. 269, Department of Statistical Science, University College London.*, 2007.
- A Philip Dawid. The Decision-Theoretic Approach to Causal Inference. *Causality: Statistical perspectives and applications*, pages 25–42, 2012.
- A Philip Dawid. Statistical Causality from a Decision-Theoretic Perspective. *Annual Review of Statistics and Its Application*, 2(1):273–303, 2015.
- A Philip Dawid. Decision-Theoretic Foundations for Statistical Causality. *Journal of Causal Inference*, 9(1):39–77, 2021.
- A Philip Dawid, Monica Musio, and Rossella Murtas. The Probability of Causation. *Law, Probability and Risk*, 16(4):163–179, 2017.
- Richard Dedekind. *Was sind und was sollen die Zahlen?* Friedrich Vieweg und Sohn, Braunschweig, 1888.
- Rick Durrett. *Probability: Theory and Examples*, volume 49. Cambridge university press, 2019.
- Kai Epstude and Neal J Roese. The Functional Theory of Counterfactual Thinking. *Personality and social psychology review*, 12(2):168–192, 2008.
- Richard P. Feynman, Robert B. Leighton, and Matthew Sands. *The Feynman Lectures on Physics*. Addison–Wesley, 1964.
- Patrick Forré and Joris M Mooij. Causal Calculus in the Presence of Cycles, Latent Confounders and Selection Bias. In *Uncertainty in Artificial Intelligence*, pages 71–80. PMLR, 2020.
- James Franklin. *The Science of Conjecture: Evidence and Probability before Pascal*. JHU Press, 2015.
- Benjamin Franta. Early Oil Industry Disinformation on Global Warming. *Environmental politics*, 30(4):663–668, 2021.
- David Galles and Judea Pearl. Axioms of Causal Relevance. *Artificial Intelligence*, 97(1-2):9–43, 1997.
- David Galles and Judea Pearl. An Axiomatic Characterization of Causal Counterfactuals. *Foundations of Science*, 3(1):151–182, 1998.
- Tobias Gerstenberg. Counterfactual Simulation in Causal Cognition. *Trends in Cognitive Sciences*, 28(10):924–936, 2024.
- Mariel K Goddu and Alison Gopnik. The Development of Human Causal Learning and Reasoning. *Nature Reviews Psychology*, 3(5):319–339, 2024.

- Nelson Goodman. The Problem of Counterfactual Conditionals. *The Journal of Philosophy*, 44(5):113–128, 1947.
- Trygve Haavelmo. The Statistical Implications of a System of Simultaneous Equations. *Econometrica, Journal of the Econometric Society*, pages 1–12, 1943.
- A. Hald. *A History of Mathematical Statistics from 1750 to 1930*. Wiley, New York, 1998.
- Anders Hald. *A History of Probability and Statistics and their Applications before 1750*. John Wiley & Sons, 2005.
- Joseph Y Halpern. Axiomatizing Causal Reasoning. *Journal of Artificial Intelligence Research*, 12:317–337, 2000.
- Joseph Y Halpern. From Causal Models to Counterfactual Structures. *The Review of Symbolic Logic*, 6(2):305–322, 2013.
- Herbert Lionel Adolphus Hart and Tony Honoré. *Causation in the Law*. OUP Oxford, 1985.
- Miguel A Hernán and James M Robins. *Causal Inference: What If*. Chapman & Hall/CRC, 2020.
- David Hilbert. *Grundlagen der Geometrie*. Teubner, Leipzig, 1899.
- Paul W Holland. Statistics and Causal Inference. *Journal of the American statistical Association*, 81(396):945–960, 1986.
- Duligur Ibeling and Thomas Icard. On Open-Universe Causal Reasoning. In *Uncertainty in Artificial Intelligence*, pages 1233–1243. PMLR, 2020.
- Duligur Ibeling and Thomas Icard. Comparing Causal Frameworks: Potential Outcomes, Structural Models, Graphs, and Abstractions. *Advances in Neural Information Processing Systems*, 36, 2023.
- Guido W Imbens. Potential Outcome and Directed Acyclic Graph Approaches to Causality: Relevance for Empirical Practice in Economics. *Journal of Economic Literature*, 58(4):1129–1179, 2020.
- Guido W Imbens and Donald B Rubin. *Causal Inference in Statistics, Social, and Biomedical Sciences*. Cambridge university press, 2015.
- Thomas Jech. *Set Theory*. Springer, 2003.
- Olav Kallenberg. *Foundations of Modern Probability*. Springer, 1997.
- Achim Klenke. *Probability Theory: A Comprehensive Course*. Springer, 2008.
- Andrey Kolmogorov. *Grundbegriffe der Wahrscheinlichkeitsrechnung*. Springer, 1933.

- Serge Lang. *Algebra*. Springer Science & Business Media, 2012.
- David Lewis. *Counterfactuals*. Blackwell Publishers, 1973.
- David Lewis. *On the Plurality of Worlds*, volume 322. Oxford Blackwell, 1986.
- Lars Lorch, Andreas Krause, and Bernhard Schölkopf. Causal Modeling with Stationary Diffusions. In *International Conference on Artificial Intelligence and Statistics*, pages 1927–1935. PMLR, 2024.
- Christopher G Lucas and Charles Kemp. An Improved Probabilistic Account of Counterfactual Reasoning. *Psychological review*, 122(4):700, 2015.
- Keith A Markus. Causal Effects and Counterfactual Conditionals: Contrasting Rubin, Lewis and Pearl. *Economics & Philosophy*, 37(3):441–461, 2021.
- Bruno Kacper Mlodozieniec, David Krueger, and Richard E Turner. Position: Probabilistic Modelling is Sufficient for Causal Inference. In *Forty-second International Conference on Machine Learning Position Paper Track*, 2025.
- Joris Mooij. Bipartite Graphical Causal Models, 2026. URL https://staff.fnwi.uva.nl/j.m.mooij/talks/2026_cambridge.pdf.
- Joris M Mooij, Dominik Janzing, and Bernhard Schölkopf. From Ordinary Differential Equations to Structural Causal Models: the Deterministic Case. In *Proceedings of the Twenty-Ninth Conference on Uncertainty in Artificial Intelligence*, pages 440–448, 2013.
- Stephen L Morgan. *Handbook of Causal Analysis for Social Research*. Springer, 2013.
- Rebecca B Morton and Kenneth C Williams. *Experimental Political Science and the Study of Causality: From Nature to the Lab*. Cambridge University Press, 2010.
- Paul Muentener and Elizabeth Bonawitz. The Development of Causal Reasoning. In Michael R. Waldmann, editor, *The Oxford Handbook of Causal Reasoning*, pages 677–698. Oxford University Press, Oxford, 2017.
- Junhyung Park and Yuqing Zhou. A Fine-grained Look at Causal Effects in Causal Spaces. *arXiv preprint arXiv:2512.11919*, 2025.
- Junhyung Park, Simon Buchholz, Bernhard Schölkopf, and Krikamol Muandet. A Measure-Theoretic Axiomatisation of Causality. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, pages 28510–28540, 2023.
- Junhyung Park, Fanny Yang, and Thomas Icard. Counterfactual Spaces. *arXiv preprint arXiv:2601.00507*, 2026.
- Giuseppe Peano. *Arithmetices principia, nova methodo exposita*. Fratres Bocca, Turin, 1889.

- Judea Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, 1988.
- Judea Pearl. Causal Diagrams for Empirical Research. *Biometrika*, pages 669–688, 1995.
- Judea Pearl. Causal Inference without Counterfactuals: Comment. *Journal of the American Statistical Association*, 95(450):428–431, 2000.
- Judea Pearl. *Causality*. Cambridge university press, 2009.
- Judea Pearl and Dana Mackenzie. *The Book of Why*. Basic Books, New York, 2018.
- Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of Causal Inference: Foundations and Learning Algorithms*. The MIT Press, 2017.
- Jonas Peters, Stefan Bauer, and Niklas Pfister. Causal Models for Dynamical Systems. In *Probabilistic and Causal Inference: The Works of Judea Pearl*, pages 671–690. ACM, 2022.
- Spencer Peters and Joseph Halpern. A Unifying Framework for Causal Modeling with Infinitely Many Variables. *Journal of Artificial Intelligence Research*, 83, 2025.
- Robert N Proctor. The History of the Discovery of the Cigarette–Lung Cancer Link: Evidentiary Traditions, Corporate Denial, Global Toll. *Tobacco control*, 21(2):87–91, 2012.
- Paul R Rosenbaum. *Observational Studies*. Springer, 2002.
- Paul R Rosenbaum. *Design of Observational Studies*, volume 2. Springer, 2020.
- Kenneth J Rothman, Sander Greenland, Timothy L Lash, et al. *Modern Epidemiology*, volume 3. Wolters Kluwer Health/Lippincott Williams & Wilkins Philadelphia, 2008.
- P Rubenstein, S Bongers, B Schölkopf, and JM Mooij. From Deterministic ODEs to Dynamic Structural Causal Models. In *34th Conference on Uncertainty in Artificial Intelligence (UAI 2018)*, pages 114–123. Curran Associates, Inc., 2018.
- Donald B Rubin. Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies. *Journal of educational Psychology*, 66(5):688, 1974.
- Donald B Rubin. Bayesian Inference for Causal Effects: The Role of Randomization. *The Annals of statistics*, pages 34–58, 1978.
- Walter Rudin. *Functional Analysis*. International Series in Pure and Applied Mathematics. McGraw-Hill, New York, 2 edition, 1991.
- Bertrand Russell. On the Notion of Cause. In *Proceedings of the Aristotelian society*, volume 13, pages 1–26. JSTOR, 1912.

- Kayvan Sadeghi and Terry Soo. Axiomatization of Interventional Probability Distributions. *Biometrika*, 112(1):asae043, 2025.
- Glenn Shafer and Vladimir Vovk. The Origins and Legacy of Kolmogorov’s Grundbegriffe. *arXiv preprint arXiv:1802.06071*, 2018.
- Oscar Sheynin. Theory of Probability. A Historical Essay. *arXiv:1802.09966*, 2018.
- Raghav Singal and George Michailidis. Axiomatic Effect Propagation in Structural Causal Models. *Journal of Machine Learning Research*, 25(52):1–71, 2024.
- Robert Stalnaker. *Ways a World Might Be: Metaphysical and Anti-Metaphysical Essays*. Oxford University Press, 2003.
- Nicole Van Hoeck, Patrick D Watson, and Aron K Barbey. Cognitive Neuroscience of Human Counterfactual Reasoning. *Frontiers in human neuroscience*, 9:420, 2015.
- Julius von Kügelgen, Abdirisak Mohamed, and Sander Beckers. Backtracking Counterfactuals. In *Conference on Causal Learning and Reasoning*, pages 177–196. PMLR, 2023.
- Michael Waldmann. *The Oxford Handbook of Causal Reasoning*. Oxford University Press, 2017.
- Naftali Weinberger. Comparing Rubin and Pearl’s Causal Modelling Frameworks: A Commentary on Markus (2021). *Economics & Philosophy*, 39(3):485–493, 2023.
- James Woodward. *Making Things Happen: A Theory of Causal Explanation*. Oxford university press, 2005.
- Sewall Wright. The Relative Importance of Heredity and Environment in Determining the Piebald Pattern of Guinea-pigs. *Proceedings of the National Academy of Sciences*, 6(6):320–332, 1920.