Causation and Explanation

Stathis Psillos

First published in 2002 by Acumen

Acumen Publishing Limited
15a Lewins Yard
East Street
Chesham
Bucks HP5 1HQ
www.acumenpublishing.co.uk

ISBN: 1-902683-41-2 (hardcover)
ISBN: 1-902683-42-0 (paperback)

British Library Cataloguing-in-Publication Data
A catalogue record for this book is available from
the British Library.

Designed and typeset by Kate Williams, Abergavenny.
Printed and bound by Biddles Ltd., Guildford and King's Lynn.

For my daughter *Demetra*

Contents

Preface

When I first started thinking about this book, I intended to write a short introduction to the philosophical debates surrounding the concepts of causation and explanation. Part of my motivation was the thought that, although there are quite a few splendid books on either causation or explanation, there was no book treating these two subjects together and aiming to cast light on their connections. As time passed, the short introduction grew bigger and bigger. In fact, it dawned on me that I couldn't adequately address the *link* between causation and explanation, unless I also wrote extensively about the laws of nature. And that's what I did. As a result, the title of the book should have been *Causation, Laws and Explanation*. In the end, the word *laws* was left out, yet the middle part of the book (Laws of Nature) forms its backbone.

The book is broad in scope, but by no means comprehensive. It aims to introduce students to the *main* theories of causation, laws and explanation. But it also ventures into more recent developments and approaches in these areas, aiming to show that, far from being philosophically sterile, these topics are very much alive and exciting. The book presupposes little knowledge of either metaphysics or philosophy of science and can be used in intermediate and advanced undergraduate courses. Yet I very much hope that professional philosophers, even specialists in these areas, will find it useful. The book presents no new theory of either causation, or laws or explanation, but it does try to evaluate, critically discuss and draw connections among existing approaches, including some very recent ones. Having the firm belief that a book aimed to be a textbook should *not*

be partisan, I have tried to be fair in my assessment of the views I discuss. Yet I make no secret of my own view that, suitably understood, the regularity view of causation, the regularity view of laws and a nomological approach to explanation are still serious philosophical contenders.

During the early stages of my work for this book, I had the privilege of extended discussions with Wes Salmon. Wes had promised me to make detailed comments on the final draft, but his untimely death deprived me of this great honour (and the philosophical community of a first-rate philosopher and a real gentleman). I am sure the book would have been much better had Wes looked at it. Another great source of inspiration (as for very many others) has been David Lewis's work. All those who work on causation, laws and explanation will feel his untimely loss. I owe a great intellectual debt to all philosophers whose views I discuss in the book, but I feel that I have to make a special reference to the inspiration I got from the work of David Armstrong, Nancy Cartwright, John Earman, Carl Hempel, Philip Kitcher, J. L. Mackie, D. H. Mellor, Peter Railton, F. P. Ramsey and Barry Stroud. Once more, David Papineau has been a great teacher for me. His encouragement has been instrumental for the completion of the book. Two anonymous readers for Acumen have made important suggestions and comments, which, I hope, have led to a substantial improvement of the final product. Their hearty endorsement of the draft they read gave me the strength I needed to carry on. Some other colleagues and friends gave me thorough comments on several portions of the book. I should particularly like to thank Helen Beebee, Alexander Bird, Mark Lange and Rebecca Schweder. The graduate students who attended a course based on this book last year were wonderful critics of both its content and its style. And my colleagues in the Department of Philosophy and History of Science at the University of Athens created a very friendly environment for me to work. The errors that, I am sure, still remain would have been many more without the generous help I received from all those people. Steven Gerrard at Acumen deserves special mention for his support throughout the completion of this book. Many thanks should also go to Jenny Roberts, who copy-edited this book, for the care with which she read the typescript.

A good part of this book was written during the summer of 2001 in Platania, a beautiful village at the foot of Mount Idi in Crete. I am

indebted to my family and other local people for making my stay there so comfortable and for giving me the space I needed to work. My wife Athena has been a source of endless care and love (as well as a great teacher of how psychologists think about causation). Finally, a couple of months after the book had been sent to the readers, our daughter, Demetra, was born. I couldn't wait to dedicate it to her.

S. P.
Athens, May 2002

Abbreviations

ADT	Armstrong-Dretske-Tooley
BT	the basic thesis: all causal explanations of singular events can be captured by the deductive–nomological model
CE	the thesis that all explanation of singular events is causal explanation
CI	causal interaction
CQ	conserved quantity
DN	deductive–nomological
DNP	deductive–nomological–probabilistic
DS	deductive–statistical
HNS	Humean nomic supervenience
HS	Humean supervenience
IS	inductive–statistical
MRL	Mill–Ramsey–Lewis
MT	mark-transmission
NT	the set of all laws
PLV	the Principle of Limited Variety
PUN	the Principle of Uniformity of Nature
RMS	the Requirement of Maximal Specificity
RVC	the Regularity View of Causation
RVL	the Regularity View of Laws
SR	statistical-relevance

Introduction

The birth of our daughter was the *cause* of great happiness to my wife and me. This *explains* why I decided to dedicate this book to her. It also *caused* certain changes in our life (for instance, that our study at home had to be converted to a nursery). It *brought about* a delay in the completion of the current book, which (hopefully) *explains* why this book might well be a bit better than it would have been had I rushed to finish it. It is *because of* her birth that I have come to realize how challenging and exciting parenthood is. And this *explains* my recent interest in books about babies. And so on and so forth.

Causal and explanatory talk is so pervasive in our everyday life, as well as in the sciences, that its importance can hardly be exaggerated. We search for causes and look for explanations in order to *understand* how and why things around us are the way they are, or behave and change in certain ways. But we also search for causes and look for explanations in order to *intervene* in the course of nature (or in the course of events, in general) and bring certain effects about or prevent others from occurring. We are interested in causation and explanation because we are *thinkers* and *agents*, because we are both theoretical and practical beings. We are worried, for instance, about the future because of certain recent developments (e.g., the destruction of the environment, or the revival of terrorism and of war, or the depletion of natural resources, or the resurgence of racism and xenophobia). We think (rightly) that we have identified at least some of the causes of this worry and we want to act to eliminate them, as well as their own causes. Theory

and practice are cemented together by the chains of causation. We offer *reasons* for our actions. But reasons, as Donald Davidson has famously stressed, can be causes of action. Besides, learning the causes of an event (be it the breaking out of the Second World War, or the collapse of a bridge, or the famine in Ethiopia) enhances our knowledge of why it happened as well as our ability to prevent similar events from happening. It also gives us *reasons* to form well-supported beliefs.

It is, of course, one thing to acknowledge, or stress, the centrality of the concepts of causation and explanation in our intellectual and practical life, and quite another thing to say what these concepts are concepts *of*. What is causation and what explanation? And how exactly are the two related? Answering these questions is the job of philosophers. Intuitively, explanation and causation go hand-in-hand. Isn't it a platitude, after all, that in order to explain something, you need to cite its causes? This platitude might not be fully general, since there are non-causal explanations (most typically, mathematical explanations), but it seems to go a long way in highlighting the *link* between causation and explanation: causes do explain and explanation does proceed via stating causes. But can we go beyond this platitude? In particular, can we offer adequate theories of causation and explanation? Can we specify the semantics of causal and of explanatory talk?

Such questions become even more urgent if we take into account that, as of late, philosophers seem to make free use of the concepts of causation and explanation in their attempts to analyse and solve traditional philosophical problems. We now have causal theories of *knowledge*, causal theories of *perception*, causal theories of *reference*, causal theories of *identity through time*, causal-role theories of *mental states* and so on. All these are piled up on top of more traditional problems such as the problem of *mental causation* (how, that is, the mental can act causally on the physical), or the problem of what exactly are the *relata* of (i.e. the things that are related by) causal relations (events, facts, states of affairs, etc.), or the problem of the *explanatory autonomy* of the special sciences, or the nature of *inference to the best explanation*. It seems either that we have to appeal to some shaky prephilosophical intuitions about causation and explanation when we investigate all of the above, or else that we need to do some serious groundwork to clarify what exactly we

refer to when we speak of causation and explanation. A central aim of this book is to present and critically discuss some of this groundwork in an attempt to clarify some of the basic conceptual issues that are involved in the philosophical debates about causation, laws of nature and explanation.

Philosophers have long disagreed about the nature of causation and explanation. They have offered different theories, either within the same or within rival metaphysical agendas. In Michael Scriven's (1975: 3) apt words, the concepts of causation and explanation "enjoy a curious love–hate relationship" with philosophers. Most think that these concepts are central to all our thinking about science (as well as about our everyday affairs) and try hard to analyse them, but there are some who *deny* their importance and suggest that the sooner they are fully dispensed with the better.

Causation

Let's start with causation. Perhaps the most famous denier of causation was Bertrand Russell (1918), who actually thought that the concept of causation was incoherent. But this was just as well for him, since, as he alleged, physics has stopped looking for causes: for "there are no such things". Here is his famous dictum: "The law of causality, I believe, like much that passes muster among philosophers, is a relic of a bygone age, surviving, like the monarchy, only because it is erroneously supposed to do no harm".[1] Now, even if Russell were right about physics – although what he asserts with confidence is, to say the least, debatable – he is definitely wrong about the other sciences. Even a cursory look at subjects such as economics, psychology and biology will persuade the non-believer that scientists do hunt for causes and assume that causes are there to be hunted. Quite powerful methods (such as randomized controlled experiments – *aka* clinical trials – and causal modelling) have been developed to aid the discovery of causal dependencies between factors and magnitudes. Far from having survived because of the erroneous supposition that it does no harm, the search for causes has been both successful and beneficial.

In a fashion similar to Russell's, Rudolf Carnap also noted that, strictly speaking, the concepts of "cause" and "effect" are meaningful only within the "perceptual world", and that, having as their

domain of application that world, "they are infected with the imprecision which attaches to concept formations within this world" (1928: 264). This may well be so, but all that it implies is that the concept of causation needs analysis and regimentation. As a matter of fact, the project initiated by Carnap and his fellow logical empiricists (but also followed by other eminent philosophers) was an attempt to *characterize* and *rescue* a legitimate core from the concept of causation, by equating causation with *de facto* invariable succession or actual regularity. This is what has come to be known as the Regularity View of Causation (RVC). It is typically seen as offering a *reductive* account of causation. As with all reductive accounts, causal talk becomes legitimate, but it does not imply the existence of a special realm of causal facts that make causal talk true, since its truth conditions are specified in non-causal terms, that is, in terms of spatiotemporal relations and actual regularities.

Most of the empiricists' ammunition has come from Hume's critique of causation. Ever since David Hume's work, philosophers of an empiricist persuasion have thought that the concept of causation is too mysterious or metaphysical to be taken seriously without any further analysis. Hence, they engaged in an attempt to *demystify* causation. They thought that the main culprit was the idea that causation implies the existence of *necessary connections* in nature, that is, connections between the causally related events that make it, somehow, necessary (or inescapable) that the effect follows from the cause. Hume was taken as the great denier of such necessary connections and as the one who conclusively showed that there were no such things to be found in nature. This denial of necessary connections in nature may be seen as the hallmark of modern Humeanism.

Some Humeans (most notably John Stuart Mill and John L. Mackie) advanced more sophisticated versions of RVC. A prominent thought has been that causation should be analysed in terms of *sufficient and necessary conditions* (roughly, an event *c* causes an event *e* if and only if (iff) there are event-types *C* and *E* such that *C* is necessary and sufficient for *E*). Another one has been that to call an event *c* the cause of an event *e* is to say that there are event-types *C* and *E* such that *C* is an *insufficient* but *necessary* part of an *unnecessary* but *sufficient* condition for *E* – *aka inus* condition). A rather important objection to Humeanism has been that regularity is not

sufficient for causation. There are too many regularities in nature and not all of them are, intuitively, causal. So Humeans have been inegalitarians towards regularities. They have tried to characterize the kind of regularity that can underpin causal relations by tying causation to laws of nature. However, other philosophers who advocate Humeanism downplay the role of regularities (or laws) in causation. A rather prominent approach has been Lewis's account of causation in terms of *counterfactual conditionals* (roughly, an event *c* causes an event *e* iff if *c* hadn't happened then *e* wouldn't have happened either). To be sure, regularities do enter the counterfactual approach to causation but in a roundabout way: as means to capture the conditions under which counterfactual assertions are true.

Many non-Humean theories deny forcefully that the analysis of causation need involve regularities, either directly or indirectly. John Curt Ducasse's *single-difference* account (roughly that an event *c* causes an event *e* iff *c* was the last – or the only – difference in *e*'s environment before *e* occurred) takes causation to link individual events independently of any regular association that there may or may not be between events like the cause and events like the effect. Salmon's *mechanistic* approach (roughly that an event *c* causes an event *e* iff there is a causal process that connects *c* and *e*) stresses that there is a local tie between a particular cause and a particular effect. Causation, non-Humeans argue, is essentially singular: a matter of *this* causing *that*.

Some philosophers think, contra Hume, that causation is directly *observable*. Others take it to be a *theoretical* relation, posited to explain a particularly robust connection between some events. Many philosophers think that if we are to avoid methodological obscurity and conceptual circularity, we have to cast the analysis of causation in non-causal terms. However, others argue that such an analysis is impossible. They dispel the charge of circularity by arguing that the concept of causation is so *basic* that it cannot be really analysed in non-causal terms. All that can be done, they claim, is offer an enlightening account of how several causal concepts are interlinked and mutually understood.

I shall now try to offer a brief diagnosis as to why there is no general agreement among philosophers on what causation is. In a nutshell, the diagnosis is that the concept of causation seems to be

characterized by conflicting intuitions, which, although almost equally central, cannot be all accommodated in a single theory. A number of philosophers – most notably D. H. Mellor (1995), Peter Menzies (1996) and David Armstrong (1999) – have recently tried to approach causation via what they have called "platitudes".[2] Perhaps, it's best to make a distinction between *platitudes* (assuming that there are some platitudinous features of causation that any theory should accommodate) and *intuitions* (assuming that there are some firm prephilosophical views about what causation is). Some of the platitudes of causation are these:

- The *difference* platitude: causes make a difference – that is, things would be different if the causes of some effects were absent. This platitude is normally cast in two ways: the *counterfactual* way – if the cause hadn't been, the effect wouldn't have been either; and the *probabilistic* way – causes raise the *chances* of their effects – that is, the probability that a certain event happens is higher if we take into account its cause than if we don't.
- The *recipe* platitude: causes are recipes for producing or preventing their effects – that is, causes are the means to produce (or prevent) certain ends (effects).[3] This platitude is normally cast in terms of *manipulability*: causes can be manipulated to bring about certain effects.
- The *explanation* platitude: causes explain their effects, but not vice versa.
- The *evidence* platitude: causes are evidence for their effects – that is, knowing that *c* causes *e*, and knowing that *c* occurred, gives us (some) reason to expect that *e* will occur.

It's not hard to agree that each and every theory of causation should accommodate these platitudes, that is, show how each of them is brought out by whatever, according to the theory, constitutes the relation of cause and effect. But there are two central intuitions about causation that also need to be taken into account.

- The *intrinsic-relation* intuition: whether or not a sequence of two distinct events *c* and *e* is causal depends wholly on the events *c* and *e* and their own properties and relations, that is, it

depends wholly on the intrinsic and local features of the actual sequence of events. For instance, according to this intuition, when we say that *the hitting with the hammer caused the smashing of the porcelain vase* what makes our assertion true has only to do with the properties of the particular hammer, the particular vase and the particular hitting.

- The *regularity* intuition: whether or not a sequence of two distinct events *c* and *e* is causal depends on whether or not events like *c* are regularly followed by events like *e*. This intuition is captured by the dictum "same cause, same effect" and is underpinned by an epistemic consideration; namely, that we are unwilling to pronounce a sequence of events *c* and *e* causal unless there has been a regular association between events like *c* and events like *e*. For instance, according to this intuition, when we say that *the hitting with the hammer caused the smashing of the porcelain vase* what makes our assertion true has to do with the fact that the hitting of porcelain vases with hammers is regularly followed by the smashing of the vases.[4]

Now, these two intuitions pull in contrary directions. The regularity intuition implies that a sequence of events is causal if and only if it instantiates a regularity. Hence, it implies that the relation of cause and effect is *extrinsic* to its relata. It makes causation dependent on *general* facts: on what happens at other places and at other times. The intrinsic-relation intuition opposes all this. It takes causation to be wholly dependent on *singular* facts: on what happens there and then, in the actual sequence of events, independently of any regularities. It would be a daunting (not to say just outright impossible) task to advance a theory that respects both of these intuitions. Most typically, Humeans base their theories on the regularity intuition, while non-Humeans base theirs on the intrinsic-relation one. A somewhat detailed investigation of the distinction between Humean and non-Humean approaches has to wait until the end of Chapter 4 (section 4.5), where, after we have examined several accounts of causation, we shall offer a map of the terrain.

It would do no harm, however, to highlight three dimensions along which the discussion about causation can be based. We have already seen the first two. The first concerns the distinction

between generalist and singularist theories. The second dimension concerns the distinction between theories that aim at an extrinsic characterization of causal relations and theories that go for an intrinsic one. The third dimension concerns the distinction between reductive approaches and non-reductive ones. Reductive approaches argue that causation is dependent on (some say it *supervenes* on) non-causal features of the world (e.g. regularities), while non-reductive accounts take causation to be ontically autonomous: an irreducible relation among events. On a first approximation, then, one could say that Humean accounts of causation take the first sides of the three distinctions: they are generalist, extrinsic and reductive. And non-Humean accounts take at least one of the remaining sides of the three distinctions: they are singularist or intrinsic or non-reductive. As further investigation will show, however, things are more complicated. Perhaps it's not very profitable to try to divide theories of causation sharply into Humean and non-Humeans, although, as we shall see, we can go some way towards achieving this task.

Laws of nature

Most Humeans have come to adopt what may be called the Regularity View of Laws (RVL): laws of nature are regularities. However, they have a hurdle to jump. For not all regularities are causal. Nor can all of them be deemed laws of nature. The night always follows the day, but it is *not* caused by the day. Nor is it a law of nature that all coins in my pocket are Euros, although it is a regularity. So the Humeans have to draw a distinction between the good regularities (those that constitute the laws of nature) and the bad ones (those that are merely accidental). Only the former, it was thought, can underpin causation and play a role in explanation. We shall see in some detail the various empiricist attempts to draw this distinction, and in particular, what I shall call the *web-of-laws* view.

According to this view, the regularities that constitute the laws of nature are those that are expressed by the axioms and theorems of an ideal deductive system of our knowledge of the world, and, in particular, of a deductive system that strikes the *best* balance between simplicity and strength. Simplicity is required because it disallows extraneous elements from the system of laws. Strength is

required because the deductive system should be as informative as possible about the laws that hold in the world. Whatever regularity is not part of this *best system* is merely accidental: it fails to be a genuine law of nature. The gist of this approach, which has been advocated by Mill, Ramsey and Lewis, is that no regularity, taken in isolation, can be deemed a law of nature. The regularities that constitute laws of nature are determined in a kind of holistic fashion by being parts of a structure. As we shall see, the *web-of-laws* view does succeed, to a large extent, in answering the question "What is a law of nature?". Yet its critics argue that it compromises the fully objective status that the laws of nature are, typically, taken to have. Why, they tend to ask, should lawhood have anything to do with how our *knowledge* of the world is organized in a deductive system? There is no doubt that the Humeans should try to dispel this charge of subjectivity. The good news, however, is that they can, to some extent at least, secure the objectivity of laws. But, as I shall argue, in order to do so they have to adopt a certain metaphysical picture; namely, that the world has an *objective* nomological structure. This structure, to be sure, will be a structure of *regularities*. Yet it may well be the case that it is objective relations among the elements of this structure, and *not* our beliefs about them, that determine what regularities are parts of this structure, and hence what regularities constitute laws of nature.

It should be noted, however, that even if the *web-of-laws* view can be deemed successful, there is a price to pay. By denying that there is any necessity in causation, the Humeans have to deny that there is any necessity in the laws of nature. Their non-Humean opponents then are quick to point out that without some appeal to a sufficiently strong concept of necessity, the distinction between laws of nature and accidental regularities will not be robust enough to support either causation or explanation. What, in short, emerges from their arguments is the view that lawhood cannot be reduced to regularity (not even to regularity-plus-something-that-distinguishes-between-laws-and-accidents). Lawhood, we are told, is a certain necessitating relation among properties (*universals*). It is noteworthy that *both* the Humeans and the advocates of the view that laws embody necessitating relations among properties agree that laws of nature are *contingent*. They do not hold in all possible worlds: they could be different, or there could be no laws at all. Yet, there has been a

growing tendency among non-Humeans to take laws of nature to be metaphysically necessary. A standard Humean retort to all these views is that, far from being enlightening, the notion of necessitation that is supposed to characterize the laws of nature (either as a contingent relation among properties or as embodying a stronger metaphysical sense of necessity) is wrong-headed and obscure.

Explanation

When it comes to the concept of explanation things may seem more promising. Here, the Logical Empiricist project of demystifying causation culminated in the attempts made by Hempel and his followers to analyse the concept of causation in terms of the concept of *explanation*. They thought that the latter could be made to be scientifically respectable by being itself analysed in terms of the concept of laws of nature and the concept of a deductive argument. The famous (to some notorious) deductive–nomological (DN) model of explanation has been a systematic attempt to subsume causation under causal explanation and to show that the latter can be fully understood and legitimized as a species of deductive argument, with one of its premises stating a universal law. In fact, the empiricist project was advanced further by enlarging the kind of arguments that can be explanations so as to include *inductive* arguments (and statistical, as opposed to universal, laws).

This reliance on laws makes it very pressing for the advocates of the DN model to have a neat way to distinguish between genuine laws of nature and accidentally true generalizations, for it is only the former that can be mentioned in legitimate explanations. The presence of an accidental generalization as a premise in a DN argument would amount to a cancellation of the nomological side of the argument. One can certainly deduce (a description of) the fact that this apple is ripe from the general statement "All apples in the fruit bowl are ripe" and the premise "this apple is in the fruit bowl". Yet this deduction hardly explains *why* this apple is ripe. Compare, however, the above with the following case. Intuitively, at least, the following argument is a perfectly legitimate explanation of the fact that Pluto describes an ellipse: Pluto is a planet and all planets move in ellipses. The difference, the advocate of the DN model would argue, is that *All apples are in the fruit bowl* is an accident, whereas

*All planets move in ellipses*is a genuine law. As a result of all this, the project of developing an adequate DN model of explanation can proceed only hand-in-hand with an attempt to characterize the genuine laws of nature.

The irony of the empiricist project is that what came out of the front door seemed to be re-entering from the window. For it seems that we cannot distinguish between good and bad explanations of some phenomena, *unless* we first distinguish between causal and non-causal explanations, or better between those explanations that reflect the causal connections between what-is-doing-the-explaining (the *explanans*) and what-is-explained (the *explanandum*) and those that do not. So it seems that we first need to sort out the concept of causation and then talk about causal explanation. If this is right, then the empiricist project outlined above gets things the wrong way around.

Yet there are plausible ways for modern empiricists to argue that, suitably understood in terms of the concept of *unification*, explanatory relations can still subsume causal relations under them. Put in a nutshell, the idea is that explanation proceeds via unification into a deductive system: a certain fact, or a regularity, is explained when a description of it is deduced within a unified deductive system. Causal relations, then, are said to mirror explanatory relations within an ideal unified deductive system. What is really interesting here is that the concept of unification can be connected with the *web-of-laws* view. Unification proceeds by minimizing the number of regularities that have to be accepted as *brute* (or as unexplained explainers). These regularities might well be accepted a the fundamental laws of nature and be captured by the axioms of an ideal deductive system that strikes the best balance between simplicity and strength. Such an ideal deductive system is none other than a unified deductive system. In line with the *web-of-laws* view, the fundamental laws of nature are the best unifiers. Yet those philosophers who resist the attempt to subsume causation under explanation point out that the foregoing view of explanation as unification will not deliver the goods. Not only is it possible that the world be disunified, but, more importantly, it seems that the foregoing view is unable to specify the conditions under which an explanation is correct. It seems, we are told, that we need to rely on the *causal structure of the world*; it is because the world has a certain causal

structure that some explanations are correct (those that capture this causal structure), while others are not. If these philosophers, notably Salmon, are right, then causal relations simply *cannot* mirror explanatory relations, even within an ideal unified system of the world. Rather, the opposite should be the case: explanatory relations, even within an ideal unified system of the world, should reflect (and capture) ontically prior causal relations.

In any case, not all philosophers agree that causal explanation should be tied to laws and have the form of an argument. Opponents of the DN model argue that explanation should rely on finding out the causes of the *explanandum*, but it need not cite laws: presenting information about the causal history of an event, or citing factors that raise the probability of an event to happen, or even stating some invariant relations among the *explanandum* and the *explanans* is taken to be enough for a good causal explanation.

The fact of the matter is that the concepts of causation, laws of nature and explanation form a quite tight web. Hardly any progress can be made in the elucidation of any of those, without engaging in the elucidation of at least some of the others. All we may then hope for is not strict analysis, but some enlightening account of their interconnections.

The menu

Although I have already hinted at the contents of this book, a more orderly presentation of the chapters to follow may help the readers orientate themselves better. Chapter 1 is about Hume and the setting up of RVC. It unravels the two projects that Hume was engaged in; namely, the analysis of causation as it is in the world and the analysis of the nature of causal inference. It culminates with a discussion of Hume's two definitions of causation. It ends with a short discussion of recent re-interpretations of Hume's views, which distance Hume from RVC, and which have given rise to the *new Hume debate*. Chapter 2 discusses Mill's elaboration of RVC, with special reference to his *methods of agreement and difference* for discovering causal laws. It also examines Ducasse's attempt to mount a major challenge to RVC and to motivate a singularist approach to causation based on Mill's method of difference. It criticizes a popular argument to the effect that causation is an observable relation and ends with

some discussion of Davidson's attempt to reconcile the Humean and the singularist approach. Chapter 3 discusses two major attempts to analyse causation in terms of *counterfactual conditionals*; namely, Mackie's and Lewis's. It also analyses Mackie's own formulation of RVC (which takes causes to be *inus* conditions). Finally, it ventures into a discussion of recent counterfactual approaches, such as Huw Price's and Menzies' human agency view and Daniel Hausman's and James Woodward's interventionist view. Chapter 4 (which concludes the part on causation) investigates theories that characterize the link between cause and effect in terms of some mechanism that *connects* them. After discussing views that argue that in the transition from the cause to the effect something *persists* or something gets *transferred*, it focuses on Salmon's early and later attempts to analyse physical causation in terms of causal processes and the transference of conserved quantities. It moves on to analyse Phil Dowe's attempt to analyse causation without an appeal to counterfactuals. It concludes with offering a rough *conceptual map* of the terrain of causal theories.

Part II of the book (on laws of nature) starts with Chapter 5, whose main aim is to critically discuss RVL. It starts with naïve versions of RVL, which simply equate laws with regularities, raises the issue of how RVL needs to be supplemented to account for the distinction between laws and accidents, and examines two major attempts towards such a supplementation: the view that the difference between laws and accidents is merely a difference in our epistemic attitudes towards them, and the much-promising Mill–Ramsey–Lewis (MRL) view, which takes laws to differ from accidents in that the regularities that are laws form a tight web. Chapter 6 focuses on non-Humean theories of lawhood. It analyses the view of Armstrong, Fred Dretske and Michael Tooley that laws are relations of contingent necessitations among properties. It strongly questions the notion of necessitation that they appeal to. It then moves on to discuss even stronger theories of lawhood, which take laws to be metaphysically necessary. It ends with a critical examination of recent arguments against the Humean view that laws supervene on non-nomic facts. Chapter 7 presents recent attempts to supersede the traditional framework of the debate on laws, by focusing more on methodological aspects of the role of laws. Among the issues that are being examined are Woodward's characteriza-

tion of laws as relations that are invariant-under-interventions, Cartwright's appeal to capacities, Lange's view on the collective stability of laws and Mellor's focus on the link between laws and natural properties. This chapter will end with a cost–benefit analysis of the major views of laws and will suggest that, on balance, RVL is still the best characterization of what a law of nature is.

The final part of the book (Part III) is on *explanation*. It starts, in Chapter 8, with the DN model of explanation. It highlights the Humean-empiricist project to deal with causation via the concept of explanation and critically discusses the counter-examples that were supposed to have uprooted the DN account. It tries to show with some precision what these counter-examples have and what they have *not* shown. It ends with an investigation of Lewis's theory of causal explanation. Chapter 9 extends the empiricist project to statistical explanation and discusses Hempel's inductive–statistical (IS) model and its problems, Salmon's statistical–relevance (SR) model and Railton's deductive–nomological–probabilistic (DNP) account. Chapter 10 extends further the empiricist project to the explanation of laws. It analyses Michael Friedman's and Philip Kitcher's models of explanatory unification. Finally, Chapter 11 engages with the issue of the connection between causation and explanation. Here, the challenge is whether the Humean-empiricist project can be completed – whether, that is, it can be shown that the explanatory relations are primary and that, somehow, the causal relations follow from them. It will be argued that Humeans can go a long way towards meeting this challenge, but that, in doing so, they have to adopt the realist view that the world has an objective structure, in which mind-independent regularities form a unified system. This insight, I will suggest, can be found in the work of Ramsey.

Absences

Three absences from this book require brief apology. I do not discuss (apart from a few brief mentions) issues related to *probabilistic causation*. We do rightly claim that, for instance, smoking causes lung cancer or that aspirin relieves headaches, even though there is no regular association (or deterministic connection) between smoking and lung cancer or taking aspirin and relief from headaches. Some philosophers think that this is already a good argument against the

view that causation is connected with invariable sequences or regularities. They then try to analyse causal claims in terms of probabilistic relations among variables, capitalizing on the intuition that causes (mostly, but not invariably) raise the probabilities of their effects. Some think that there are good empirical reasons to jettison determinism (roughly, the view that each and every event has a fully sufficient set of causes) in favour of indeterminism (roughly, the view that there are genuinely chancy events). They then try to show that indeterminism and causation mix well, given the thought that a certain event can be *caused* to happen even though its cause made only a difference to its chance of happening. Interestingly, these ideas are extended to deterministic causation as well, with the prime thought being that an effect is deterministically caused to happen if its probability, given its cause, is unity. It is also noteworthy that probabilistic theories of causation are advanced by both Humeans (who think that causal connections are reducible to relations of probabilistic dependence) and non-Humeans (who think that causal relations are *not* reducible to probabilistic relations but, nonetheless, take the latter to illuminate causation.) Discussing these intricate matters would have made this book unmanageably long. So the reader is advised to look at Patrick Suppes (1984), David Papineau (1985) and Ellery Eells (1991) for excellent accounts of probabilistic causation. For what it is worth, my own view is close to Hausman's (1998: 186). I too think that acceptance of indeterminism implies the acceptance of uncaused things, but that there can be fully deterministic causation of probabilistic states.

Another issue I do not discuss (apart from a few scattered observations and remarks) concerns the *direction of causation*. Why is it the case that causes precede their effects in time? Some philosophers (including Hume) thought that this feature is conceptually constitutive of causation, while others think that it is an empirical feature of the actual world, which needn't obtain in other possible worlds. Other philosophers try to define the order of causation independently of the concept of time, so that they can then explain the direction of time in terms of the direction of causation. All philosophers who have thought hard about causation have dealt with this issue of *causal priority*. But, here again, I would advise the interested reader to look at Paul Horwich (1987) and Hausman (1998) for excellent guidance into all this.

Finally, a third issue I do not touch at all relates to the so-called *pragmatics of explanation*. Some philosophers focus on the *act* or the *process* of explaining, instead of the *product* of explanation. They argue that an explanation should be seen as an answer to a why-question and note that the *relevant* answers will depend on the presuppositions or the interests of the questioner, on the space of alternatives, and, in general, on the context of the why-question. Here is Alan Garfinkel's (1981: 21) famous example. A priest asked Willie Sutton, when he was in prison, "Why did you rob banks?", to which Sutton replied, "Well, that's where the money is". Garfinkel's thought is that this is a perfectly legitimate answer for Sutton, because for him the space of relative alternatives (the contrast class) concerns robbing groceries or diners or petrol stations, and so on. But the space of relevant alternatives for the priest is quite different: not robbing anything, being honest, and so on. The difference of perspective can be brought out by placing the emphasis on different parts of the question: "Why did you rob *banks*?" as opposed to "why did you *rob* banks?" Pragmatic theories of explanation, very different in their details but quite similar in their overall focus on the act of explaining and the contrast classes, have also been offered by Bas van Fraassen (1980) and Peter Achinstein (1983).

With all this in mind, it's now time to leave the starter and move on to the main three-course meal. I hope you enjoy it.