

The background of the cover is a dark, deep blue. Overlaid on this are intricate, ethereal patterns of lighter blue and white. These patterns resemble smoke, mist, or perhaps complex scientific data visualizations like fractals or fluid dynamics simulations. They swirl and flow across the entire surface, creating a sense of movement and depth. The patterns are most concentrated in the upper left and lower right, with more subtle wisps throughout the rest of the frame.

OXFORD

CAUSATION  
IN SCIENCE  
AND THE METHODS OF  
SCIENTIFIC DISCOVERY

RANI LILL ANJUM AND STEPHEN MUMFORD

# Causation in Science and the Methods of Scientific Discovery



# Causation in Science and the Methods of Scientific Discovery

Rani Lill Anjum and Stephen Mumford

**OXFORD**  
UNIVERSITY PRESS

# OXFORD

UNIVERSITY PRESS

Great Clarendon Street, Oxford, OX2 6DP,  
United Kingdom

Oxford University Press is a department of the University of Oxford.  
It furthers the University's objective of excellence in research, scholarship,  
and education by publishing worldwide. Oxford is a registered trade mark of  
Oxford University Press in the UK and in certain other countries

© Rani Lill Anjum and Stephen Mumford 2018

The moral rights of the authors have been asserted

First Edition published in 2018

Impression:1

All rights reserved. No part of this publication may be reproduced, stored in  
a retrieval system, or transmitted, in any form or by any means, without the  
prior permission in writing of Oxford University Press, or as expressly permitted  
by law, by licence or under terms agreed with the appropriate reprographics  
rights organization. Enquiries concerning reproduction outside the scope of the  
above should be sent to the Rights Department, Oxford University Press, at the  
address above

You must not circulate this work in any other form  
and you must impose this same condition on any acquirer

Published in the United States of America by Oxford University Press  
198 Madison Avenue, New York, NY 10016, United States of America

British Library Cataloguing in Publication Data  
Data available

Library of Congress Control Number: 2018939482

ISBN 978-0-19-873366-9

Printed and bound by  
CPI Group (UK) Ltd, Croydon, CR0 4YY

Links to third party websites are provided by Oxford in good faith and  
for information only. Oxford disclaims any responsibility for the materials  
contained in any third party website referenced in this work.

# Preface

There is no doubt that there is causation in science. Some of the chief goals of science are understanding, explanation, prediction, and technical application. Only if the world has some significant degree of constancy, in what follows from what, can these scientific activities be conducted with any purpose. This regularity in nature is the basis of prediction, for example, which would otherwise be entirely unreliable. But what is the source of such constancy and how does it operate? How is it ensured that one sort of event tends to follow another? This is a question that goes beyond science itself—beyond the empirical data—and inevitably requires a philosophical approach: a metaphysical one. We argue in this book that causation is the main foundation upon which science is based. It is causal connections that ground regularity. They are the reason that one kind of thing typically follows another with enough reliability to make prediction and explanation worthwhile pursuits. It is causation that grounds the possibility of knowledge.

Should scientists concern themselves with what philosophers have to say? On this issue we argue for a resounding affirmative. Philosophy tells us about the nature of causation itself and thus what we should be looking for when we investigate the world. All scientists make presuppositions. They rest their endeavours upon assumptions that are not themselves empirically demonstrable through data. These assumptions are metascientific. Scientists conduct experiments, for instance: trials and interventions that change one thing in order to see what changes with it. Without a belief in causal connections, or in something being causally responsible for those changes, such experimentation would have no point because the intervention would not be productive of the result; also there would be no possible future application. Furthermore, the methods adopted or preferred by practising scientists must reveal something of what they take causation to be. At the very least, they must be looking for something that they accept as evidence of such causal connections.

For some time, philosophers have paid an interest in science, attempting to make their philosophical theories better informed empirically. Too often neglected is the opposite goal: of making scientific practice better informed philosophically. In the case of causation this is absolutely crucial for it can make a significant difference to how science is conducted. If one is persuaded that causes are difference-makers, for instance, this is what one will seek to find. If one thinks causes are regularities, one will instead look for them.

Now why should any of this matter to anything? Actually, it matters a lot. For example, one of the greatest challenges facing humanity is climate change. For all we know, this is a threat to our very existence, or at least that of a significant number of the world's population. It is hard to think of a greater problem than that. Rightly, we

look to science to help us sort out the facts, many of which remain controversial. There are a number of questions we can pose, all of which should have a scientific answer. The first is simply whether climate change is real, and almost all impartial scientists now agree that it is. Other questions have then become more pressing. Did human activity produce this climate change? What were the mechanisms? And, looking forward, what are likely to be the medium- and longer-term effects of such change? Will the polar ice caps all melt? Will sea levels rise? Will the Gulf Stream switch off? Will there be crop failures? Will we starve? Can it be stopped? Can it be mitigated? All of these are causal questions. They concern what-will-cause-what and we want scientific ways of uncovering these causal facts. One of our biggest ever problems has causation at its heart, therefore: causal matters that we look to science to explain and, hopefully, control.

The aim of this book is not to offer an overview of the philosophy of causation. Other books already do this. The recent *Causality* by Illari and Russo (2014) does an excellent job of presenting an impartial survey of the available theories of causation and how they relate to science. This has freed us up for a more opinionated investigation of the topic. We think that there are certain dominant philosophical views of the nature of causation that continue to inform and influence methods, or at least influence what we think the norms of scientific discovery should be. As we hope to show, however, some of those perceived norms create a conflict when they confront the scientist in everyday practice. Scientists know that they cannot perfectly preserve a norm of objectivity, that they cannot find exceptionless regularities that are indicative of causal connections, that they cannot attain absolute certainty, that causal knowledge is defeasible, and so on. Scientific practice thus provides a hard test for our received norms and philosophical theories of causation and, we argue, gives us good grounds to overthrow and replace some of those norms. We strive, therefore, for a scientifically informed philosophy as well as a philosophically informed science.

With twenty-eight chapters, it would certainly be possible for a busy reader to dip in and out of the book. The chapters can be read in any order and in isolation. But a thesis is developed in the whole that would reward a sequential reading. If you like, you could proceed straight to the conclusion. There, we list nine new proposed norms of science that are supported by the arguments in the book. Under each norm, we list which chapters are relevant: those in which support for adoption of a new norm is justified. You could then look at the relevant chapters. Another way to read the book is to skim the keywords that we have given for each chapter, instead of an analytic table of contents, and then consult the topics that interest you the most. Although we have made it easy to read the book in these ways, we still hope that some readers will choose to follow the whole text from start to finish.

Examples are used to illustrate the philosophical points. They are drawn from a range of sciences and are kept as simple as they can be. For us, it is important that readers with different backgrounds can understand and relate the discussions to their own discipline and perhaps consider more detailed and relevant examples themselves.

The book consists of eight parts. Part I contains three chapters about the philosophical motivation for the book. Chapter 1, *Metascience and Better Science*, urges that philosophers and scientists need to address the issue of causation together, aiming towards a reflective equilibrium satisfactory both theoretically and empirically. Since the scientific methods for discovering causes reveal ontological commitments as to the nature of causation, these commitments must be made explicit and subject to critical reflection. In Chapter 2, *Do We Need Causation in Science?*, we argue that causation occupies a foundational place, underpinning the purposes of science to explain, understand, predict, and intervene. Chapter 3, *Evidence of Causation Is Not Causation*, draws a distinction between how we learn about causation (epistemology) and what we take causation to be (ontology). The rest of the book is an attempt to show that these two, although they must not be confused, are tightly and inevitably linked. How we seek to establish causation must be informed by what we think causation is.

The next four chapters, in Part II, discuss an orthodox view that causation is conceptually and epistemologically linked to perfect correlations. Chapter 4, *What's in a Correlation?*, concerns how we separate causal from accidental correlations, while neo-Humeanism struggles to make ontological sense of such a distinction. Chapter 5, *Same Cause, Same Effect*, questions the theoretical assumption that causation should be robust across all contexts, since this is not supported empirically. Chapter 6, *Under Ideal Conditions*, takes this discussion further, showing how our theoretical expectation of causal necessitation is philosophically salvaged by stipulating a set of ideal conditions. In Chapter 7, *One Effect, One Cause?*, we warn against oversimplifying causes. Typically there are multiple causes of an effect and treating them in isolation can let us miss the importance of their interaction.

Part III consists of three chapters that present an alternative approach to the Humean orthodoxy. Chapter 8, *Have Your Cause and Beat It*, introduces the notion of additive interference, where an effect is counteracted by the adding of something. This explains why causation is essentially sensitive to context. In Chapter 9, *From Regularities to Tendencies*, a case is made for understanding causes as tendencies rather than constant conjunctions, conceptually detaching the notion of causation from perfect regularity. In Chapter 10, *The Modality of Causation*, we argue that causation involves a primitive and dispositional modality, weaker than necessity but stronger than pure contingency.

Part IV promotes causal theories and mechanisms as an alternative to finding causation in regularity and repetition. In Chapter 11, *Is the Business of Science to Construct Theories?*, we argue that a causal theory is needed in addition to the data if we want something more than merely mapping the facts. Chapter 12, *Are More Data Better?*, makes a case for causal singularism, where causation happens in the concrete particular, over Humeanism, where the single instance is derived from a universal claim. In Chapter 13, *The Explanatory Power of Mechanisms*, we point to how we need qualitative and mechanistic knowledge for deep causal understanding and



explanation. We here present our distinctive account of mechanisms. However, in Chapter 14, *Digging Deeper to Find the Real Causes?*, we scrutinize the reductive project of finding causal mechanisms always at lower levels of nature. We conclude that mechanisms can be higher level, and we give accounts of holism and emergence.

Part V concerns how causes are linked to effects. In Chapter 15, *Making a Difference*, we discuss the counterfactual theory and argue that it fails to account for some instances of causation. This shows that causation is not the same as difference-making. Chapter 16, *Making Nothing Happen*, develops the point that there are cases of causation where no change or event follows, but that these are some of the most important causal situations. We then look at the matter of causal chains in Chapter 17, *It All Started with a Big Bang*, and consider whether their existence shows that causation is transitive. The part ends with Chapter 18, *Does Science Need Laws of Nature?*, where we question the ontological need for universal laws in addition to intrinsic propensities and their causal interactions.

Part VI is about probability. Chapter 19, *Uncertainty, Certainty, and Beyond*, focuses on degrees of belief, such as doubt and certainty. Notably, we make room for a class of cases in which you can still be certain of something even with less evidence than before. Chapter 20, *What Probabilistic Causation Should Be*, is where we offer our theory of worldly chance, based on a distinctive account of propensities. In Chapter 21, *Calculating Conditional Probability?*, we show that our intuitive notion of conditional probability must be separated from the standard ratio analysis of  $P(A|B)$ , for epistemological and ontological reasons.

Part VII is primarily on the problem of external validity and contains two chapters: Chapter 22, *Risky Predictions*, and Chapter 23, *What RCTs Do Not Show*. The first of these argues that our causal predictions are essentially fallible but useful nevertheless. Indeed, any theory of prediction that did not account for its fallibility would be flawed because of that. The second points to some significant shortcomings of these large-scale population studies in dealing with causal factors such as individual variation, heterogeneity, complexity, and marginal groups.

Part VIII contains five chapters on causal discovery. Chapter 24, *Getting Involved*, develops the manipulationist insight, but mainly to argue that our causal knowledge happens in close interaction with the world, not by distanced observation. Chapter 25, *Uncovering Causal Powers*, presents our account of technological innovation, which we argue rests mainly in teasing out the hidden powers of things. Chapter 26, *Learning from Causal Failure* is a consideration of the opportunities of new knowledge that arise from unsuccessful experiments and discrepancies. Given that there is a diminishing return in confirming evidence, after a point, big breakthroughs are more likely to follow from negative results. In Chapter 27, *Plural Methods, One Causation*, we argue for a combination of epistemic pluralism and an ontological monism. The failure of analysis has led others to assume that causation in reality must be many different things. But this overlooks the other possibility: causation is one thing but primitive. We must then investigate it through its

symptoms and our methods must detect those. In Chapter 28, *Getting Real about the Ideals of Science*, we challenge the motivation for idealization and abstraction in science when dealing practically with a more messy reality. There is a crisis in reproducibility, which shows how some of our expectations of science are unrealistic and based on a mistaken notion of causation.

What emerges from all these considerations is that, if we understand causation right, we should approach it in a certain way through our scientific methods. We finish by offering nine new norms for causal science, drawing together the many themes of the book.

This book is the culmination of the research carried out on the *Causation in Science* (CauSci) project at the Norwegian University of Life Sciences (NMBU) from 2011 to 2015, funded by the FRIPRO scheme of the Research Council of Norway. We thank the CauSci team (original and substitute), as well as the many project collaborators, for engaging with our work. The material was further developed through a teaching course at NMBU and we are grateful for the discussions and input from the students of PHI302 and PHI403. We have also received invaluable contributions from the CauseHealth team and collaborators. In particular, we want to thank Fredrik Andersen for raising and discussing with us many of the issues of the book. We are grateful to Johan Arnt Myrstad, the *sine qua non* of Chapter 21. A special thank you to Elena Rocca, our practising scientist in residence, for suggesting and researching many of the examples we use.



# Contents

<i>List of Figures</i>	xv
------------------------	----

## Part I. Science and Philosophy

1. Metascience and Better Science	3
<i>science – philosophy – norms</i>	
2. Do We Need Causation in Science?	11
<i>scepticism – observation – empiricism</i>	
3. Evidence of Causation Is Not Causation	19
<i>ontology – epistemology – operationalism</i>	

## Part II. Perfect Correlation

4. What's in a Correlation?	29
<i>regularities – statistics – coincidences</i>	
5. Same Cause, Same Effect	37
<i>robustness – irregularities – falsification</i>	
6. Under Ideal Conditions	46
<i>idealization – generality – exception</i>	
7. One Effect, One Cause?	53
<i>complexity – background conditions – triggers</i>	

## Part III. Interference and Prevention

8. Have Your Cause and Beat It	63
<i>interference – intervention – causal expansion</i>	
9. From Regularities to Tendencies	71
<i>incidence – tendencies – vector models</i>	
10. The Modality of Causation	79
<i>dispositionality – necessity – powers</i>	

## Part IV. Causal Mechanisms

11. Is the Business of Science to Construct Theories?	89
<i>raw data – theory development – underdetermination</i>	

12. Are More Data Better?	97
<i>induction – causal uniqueness – dispositionalism – evidence base</i>	
13. The Explanatory Power of Mechanisms	104
<i>mechanisms – quantitative – qualitative</i>	
14. Digging Deeper to Find the Real Causes?	112
<i>reductionism – levels – emergence – demergence</i>	

## Part V. Linking Causes to Effects

15. Making a Difference	123
<i>comparison – counterfactuals – overdetermination</i>	
16. Making Nothing Happen	131
<i>stability – homeostasis – absences</i>	
17. It All Started with a Big Bang	139
<i>causal chains – determinism – transitivity</i>	
18. Does Science Need Laws of Nature?	146
<i>governance – properties – covering law</i>	

## Part VI. Probability

19. Uncertainty, Certainty, and Beyond	157
<i>probability – Bayesianism – credence</i>	
20. What Probabilistic Causation Should Be	165
<i>chance – frequentism – propensities – overdisposing</i>	
21. Calculating Conditional Probability?	174
<i>ratio formula – conditionals – compositionality</i>	

## Part VII. External Validity

22. Risky Predictions	185
<i>uncertainty – fallibility – external validity</i>	
23. What RCTs Do Not Show	194
<i>variation – exclusion – bias – policy</i>	

## Part VIII. Discovering Causes and Understanding Them

24. Getting Involved	205
<i>manipulation – presuppositions – processed data</i>	
25. Uncovering Causal Powers	212
<i>technology – hidden powers – side effects</i>	

26. Learning from Causal Failure	220
<i>diminishing returns – negative results – discrepancies</i>	
27. Plural Methods, One Causation	227
<i>pluralism – primitivism – evidence hierarchies</i>	
28. Getting Real about the Ideals of Science	237
<i>reproducibility – replicability – messiness</i>	
Conclusion: New Norms of Science	245
<i>References</i>	255
<i>Index</i>	271



# List of Figures

5.1. Model versus data	40
5.2. An outlier	42
9.1. A vector model of two tendencies	74
9.2. Additive interference	74
9.3. Subtractive interference	74
20.1. Asymptotic relation between extent of propensity and probability	171
22.1. A vector with the resultant tendency R	189
27.1. Causation and its symptoms	232





# PART I

## Science and Philosophy



# 1

## Metascience and Better Science

### 1.1 What Science Is and Should Be

Science is many things. We might say that it is an activity aimed at discovering a class of important truths about the world; and then showing what can be done with that knowledge. We can think of science as a tool, a method, or, more generally, a philosophy. Central to science, perhaps even constitutive of it, is a set of norms for the correct, systematic acquisition of empirical knowledge. By a norm, we mean a standard of what ought, or ought not, to be done. For scientific norms, the concern is what we ought to do or ought not to do in order to best gather and utilize knowledge about the world. Empirical knowledge is knowledge derived from experience, as opposed to mathematical, logical, and conceptual truths, which can be discovered through use of reason.

For example, it is arguable that for knowledge to count as scientific, it needs to be objective. It cannot be just one person's view, for instance. Hence, if you claim to have found something remarkable under your bed that no one else is allowed to see, the report, in that state, will not count as scientific and will thus not be an admissible datum for any scientific theory. Being just one person's experience, this report will be classed as subjective and unverified. However, if knowledge is acquired according to the norms of science, including the norm that it be objective, then it may qualify for the status of scientific. The norms are thus important to many forms of enquiry because it often matters that a claim is scientific. It means it should be trustworthy, for instance. A justification of science could be that it is knowledge acquired through these norms, which have been accepted as the right way to know the world. Science is valuable, it might be said, because it is objective, among other things.

There are a number of norms for correct empirical discovery. It needs to be admitted, however, that these norms continue to be contested. The history of science shows multiple, ongoing debates about what is the correct scientific method. Should science start with the recording of data, for instance, or is it ever permissible—indeed inevitable—to begin with a hypothesis that you can then subject to testing? There is still worthwhile discussion to be had on this question. It may now sound surprising to hear some of the older debates, especially on subjects we think are clear-cut. Galileo (1632) offered an experimental method for science, for example; but experimentation was controversial at the time (Gower 1997: 23). We now think it obvious that science

should conduct experiments. In a tradition dating back to Plato (see the allegory of the cave, *Republic* 514a–520a), however, the senses were regarded as unreliable so we were urged to trust reasoning instead, which was precise and conclusive. Even in Galileo's time, evidence from the largely unaided senses was imprecise and uncertain, so thought to have only limited use for science. Galileo compromised. Most of the experiments he proposed were thought experiments. He reasoned through what he thought had to happen, including, it seems, in the famous experiment in which two objects of the same size but different weights were thrown from a tall building. It was the use of reason that led Galileo to conclude that they would fall at the same rate. The experiment could in principle be performed 'for real' but there seemed no cause to do so, thought Galileo, since any result other than his own would be contrary to reason and therefore could not happen (Gower 1997: 30ff).

We might now think that we possess accurate measuring devices within experimental set-ups and that we can trust our observations, which are as precise as pure reasoning. But this is still contestable. Sometimes there is such confidence in the reasoning within a theory that it hardly needs any empirical confirmation. Or, as in theoretical physics, we are not yet able to empirically confirm the theory and thus have to rest the argument on purely rational considerations. Opposing this view is the thought that reason itself can no longer be trusted, in that the world needn't work in a particularly rational or intuitive way. Some results in contemporary physics seem to offend common sense, for instance. In one well-known instance, we are asked to accept that Schrödinger's cat is both dead and alive at the same time and that this explains how some aspects of the world really work. We might have to revise our view of what is logical.

We said that the norms of science were many. It may help if we offer up some initial candidate norms. All of them have an initial plausibility and they might collectively be taken as a reasonable default or starting position. We could say, then, that the following list captures much of what science ought to do:

- Be objective. For example, results should be reproducible so that others can observe the same experiment and record the same outcome.
- A theory should 'preserve the data', i.e. be consistent with it eventually, more or less.
- The more empirical evidence there is in favour of a theory, the more likely it is to be true and thus the more inclined one should be to accept it.
- Favour theories that have greater explanatory power or scope. If a new theory can explain everything that the old theory explained, and more, that is a reason to favour the new theory over the old one.
- Predictive success counts in favour of a theory in that one sign of a good or true theory is that it gets confirmed by subsequent results.

A further reason for calling these norms 'default' is that we want to leave open the question of whether they are strict, 'hard and fast', or whether they should be treated

as heuristic ‘rules of thumb’, admitting exceptions. After all, perhaps there are some cases where you should favour a theory with narrower scope over one with broad scope, such as if the broader theory rests on implausible assumptions. We should allow that the default norms may be revised or given more precision. Predictive success is not a straightforward matter, for instance. There are problems such as confirmation bias and theory-dependence of observation where it could look like a theory has achieved predictive success but hasn’t really. One would have to rule out such cases in order to preserve the integrity of the norm. If one can’t do that, one might have to abandon the norm.

## 1.2 Why Philosophy?

This account shows that there are at least two jobs for the philosophy of science. First, we have to decide what the correct norms of science are and justify them. Is it right to favour the theory with the most empirical evidence, for instance; and if so, why? Second, once all the norms are in place, we must also consider what the best ways are of satisfying them. If, for example, we agree that science should be objective, one then has to consider whether this demands reproducibility of results (see Chapter 28). Does it mean that anyone, anywhere, should be able to follow the method and get the same outcome as the one reported? Are there other ways of satisfying the requirement of objectivity? Are any of them more important than the norm of reproducibility?

Why, it might be wondered, do we say that these questions are for philosophy to answer, rather than science itself? Can’t science work it out unaided? Isn’t it presumptuous to claim that science needs philosophy? There is indeed a view that philosophy is of little or no relevance to this discussion, nor is anything else outside of science. Scientism, sometimes called naturalism, is the view that every meaningful question can be answered scientifically. Insofar as philosophy is attempting a task which is non-empirical, hence non-scientific, scientism questions philosophy’s right to pronounce about the nature of the world. Some proponents of this view are within science. Stephen Hawking (2011: 5) has declared that philosophy is dead because physics has answered all its substantial questions. But some defences of scientism also come from within philosophy, such as from Ladyman et al. (2007), who are following a naturalist tradition promoted by Quine (1995, for one instance). We wish to overcome any such conflict between philosophy and science. Being pro-philosophy does not mean one is anti-science. Similarly, being pro-science should not mean being anti-philosophy. Both are needed, as we aim now to explain.

First, if our task is to understand what the correct norms of science are, the answer clearly cannot come from within science itself, for any such answer would be question begging. One might be tempted by the opposite view. When asked how science should be done, someone could say that one should look at science and see how it is done, or that one should consult a trained scientist who knows exactly. But this is no good. We want to know how science ought to be conducted, which isn’t the same as how it is

actually conducted. If we take the norms to be definitive or constitutive of science, then clearly one can only consider the validity of those norms by stepping outside of the scientific practice itself. Consideration of the normative aspect of science is thus inherently a philosophical and abstract enterprise rather than an empirical one. The key issue is whether the current norms of a practice could be wrong or incomplete. One cannot make this judgement from within that practice.

Second, it has to be acknowledged that there is philosophy in science whether one likes it or not. Science rests on philosophical assumptions, including metaphysical ones. These assumptions cannot be proven by science itself, but only assumed, and this shows us that scientism is untenable. This is effectively admitted during an attempted defence of scientism by Ladyman et al.: ‘With respect to Lowe’s... claim [Lowe (2002: 6), that “Naturalism depends upon metaphysical assumptions”], it is enough to point out that even if naturalism depends on metaphysical assumptions, the naturalist can argue that the metaphysical assumptions in question are vindicated by the success of science’ (Ladyman et al. 2007: 7). This concession is decisive. We can see a number of significant points in it that are worth detailing.

First, the concession accepts the possibility that science doesn’t tell us everything. It can rest upon non-scientific assumptions. It is clear that these assumptions cannot be evaluated by science itself because they must be assumed in order to have science in the first place.

Next, the concession admits the basis of a strong argument in favour of a certain metaphysical assumption or set of assumptions. It tells us that such assumptions would be justified when science rests upon them and science is a success. We will see in Chapter 2 that this gives us a good argument for the reality of causation. Science assumes the validity of both observation and intervention and there could be neither without causation. We have good grounds, then, for saying that causation is real because science rests upon it and is successful. So this is an effective argument in favour of a foundational metaphysics, grounding a successful pursuit such as science.

Another significant point is that the Ladyman et al. concession assumes a type of justification for science that cannot come from science itself but, as argued above, is a norm that can only be evaluated from outside. This is the claim that success—for instance, explanatory, predictive, and technological—is what vindicates science. Science itself cannot show that you ought to follow the practice that produces success of this kind. It depends on what you want. Different activities are successful on different grounds: grounds that have to be chosen through rational enquiry. We are entitled to ask specifically what the markers of success are in science. What exactly is success in this case?

Finally, it is worth considering how this norm of science would be assessed. Suppose the naturalist were to claim, without resorting to philosophy, that it’s simply obvious that predictive success is the yardstick against which we should measure science. That may be so. Isn’t it just intuitive that we should justify the worth of science, and therefore its metaphysical basis, on the grounds of its success? What

more than explanatory, predictive, and technological success could one sensibly want? However, this is a defence that Ladyman et al. have already ruled out. They argue that ‘as naturalists we are not concerned with preserving intuitions at all’ (Ladyman et al. 2007: 12) and they go on to point out how many things in science are true but counterintuitive. One should not accept the argument (that science is validated by its success) on intuitive grounds, therefore; at least not according to defenders of naturalism.

Yet, this type of argument would not be quite so bad otherwise. Ladyman et al. cannot accept it because it undermines their defence of scientism. For someone who is not seeking to defend scientism, it could however be accepted. Indeed, so compelling is the intuition, that the authors use it very quickly as a basis for accepting another norm of science. To justify that science is seeking a ‘relatively unified picture’ (p. 27), they argue that the opposite would be a ‘mystery’. That sounds fair enough, except for the fact that the authors had just claimed that being intuitive is no virtue in science. What, then, can they maintain is bad about something being a mystery?

### 1.3 Philosophical Assumptions in Science

Philosophical assumptions are inevitable in science. There are then two types of scientists: those who are aware of science’s philosophical underpinnings and those who are not. Only if we are aware of what we have assumed are we able to reflect critically upon its various aspects and ask whether it really gives us a sound foundation. Without such awareness, our scientific conclusions might contain erroneous presuppositions that do not serve us well or withstand rational scrutiny.

Let us look at one way in which this might happen and, in doing so, move towards the central topic that occupies us in this book: the scientific discovery of causation. We should also understand some of the differences between philosophical and scientific approaches to knowledge of causation. While science is empirical and deals with the concrete, such as particular matters of fact, recorded as data, philosophy deals with the general and abstract. Thus, in science, including social science, we might want to find out what causes what: whether a particular genotype makes a person susceptible to cancer; whether heat increases or decreases the tensile strength of a certain metal, or whether a reading programme improves the rate of social mobility. These are all particular causal hypotheses, even though none of them uses the term ‘cause’ explicitly. They could easily have done so (‘does reading cause an increase in social mobility’, and so on). Philosophy—in particular metaphysics, which studies what there is in the most general terms (Mumford 2012)—will look at what is common to these three scientific instances. They all claim that some A is causing some B. This shows that philosophy is abstract in the sense that its concern is with the nature of causation, irrespective of what the causes and effects are. Philosophy thus studies what it is for one thing to cause another, abstracted away from the particular things that are causing and being caused.



A naturalist might claim that we can just let science tell us what it is for one thing to cause another, but that would be a mistake. Science cannot do this without ceasing to be science. To do so is to go beyond the data and enter into philosophy. Science can propose theories of what causes what, based on the data, but it exceeds the norms of science to comment on the largely non-empirical matter of what causation is itself.

Naturalists find themselves in a trap, then. One can only say that A causes B if one has an account of what it is for one thing to cause another; but science unaided by philosophy cannot judge on that matter. For this reason, some scientists decline ever to make causal claims. Causal claims are rare within a certain tradition in epidemiology, for instance, even though such a science would seem pointless unless it is discovering causes (Broadbent 2013). It might be thought that the scientific task is complete once the data are recorded: perhaps a raised incidence of some disease upon increase of another variable. A naturalist reaction to the trap is, then, to proclaim that there is no causation in science, or none in one particular science. This question is the topic of Chapter 2. We can say here, however, that this sort of claim must be based on an assumption of what causation is and is not (Kerry et al. 2012, Andersen et al. 2018). Based on that assumed view of causation, one may just conclude against its presence. But what if the nature of causation is not as the naturalist assumed? Perhaps no regularities or determinacies are found in physics, but that only counts against the presence of causation if causation has to be deterministic and issue in regularities. The problem here is that we can only find the right or best account of the nature of causation and all its features by doing at least some philosophy. The naturalist is trapped by a refusal to do any substantial metaphysics, which hinders progress.

There are of course many other notions in science that benefit from a higher-level understanding. What is an explanation, for instance, a law of nature, or a natural kind? Only with a mature understanding, of what these things are supposed to be, can one accept or reject them in science. One could not say, for instance, that there are no natural kinds in reality because nothing has feature N, if feature N is not really a part of what it is to be a natural kind. And this latter question is at least in part—and perhaps primarily—a philosophical matter. Similarly, there are many other assumptions at play in science that are not empirically grounded, concerning, for instance, the notion of complexity (see Rocca and Andersen 2017). Some accept reductionism, for example. Perhaps one science is reducible to another, such as if we can explain psychology in terms of evolutionary biology or economics in terms of neuroscience (see Chapter 14). Perhaps everything is reducible to physics. Again, though, we need to understand what conditions would qualify as a reduction, which then gives us a grip on what, if anything, would allow us to recognize a case of it in reality.

Our job, however, concerns specifically causation in science. One task is to tease out the assumptions that are being made about the nature of causation: articulate them so that we can scrutinize them and judge whether they are right. But this is not our only task; indeed it is not our main task. This book is primarily about the nature of causal science and its own philosophical presuppositions, as manifested in the

norms of causal science. What we mean by causal science is the science of finding causes in the natural world and using that knowledge. We believe that smoking tobacco causes cancer, for example, and that the heating of iron rods causes them to expand. Do we have good grounds for these beliefs? We can see that this comes down to the question of whether the correct norms of causal discovery have been followed and given a positive judgement. But do we know that those correct norms have been yet found? Furthermore, is it possible that some norms of causal discovery are inappropriate? This could again be a partially philosophical matter. Suppose one has a regularity theory of causation: that it consists in an absolute regularity or constant conjunction (Hume 1739: I, iii, 2). One may then adopt as a norm of causal science that one should record all the constant conjunctions and draw causal inferences from them. But this norm is inappropriate, we allege, because it will fail to reliably identify the causal facts. It could pronounce that smoking is not a cause of cancer, because not all smokers get cancer; or, contrariwise, that humanity is the cause of mortality, which we would say it is not, even though all humans are indeed mortal. What these examples illustrate is simply that our philosophical view of the nature of causation shapes the norms that we adopt for causal science; that is, it determines what we think we should look for when seeking causal connections.

## 1.4 The Critical Friend

The defence of a philosophical contribution to science should in no way be understood as an attack on science. Although science cannot answer every question, it still gives us the best way to answer questions of a certain kind, including the empirical matter of what is causally connected to what. Philosophy should be seen as science's critical friend, pointing out those normative aspects that in part constitute science and examining the credibility of the underlying metaphysics upon which it rests. In return, science can be philosophy's critical friend, too, allowing the abstract theory to be informed by a better understanding of the concrete reality. It might be assumed in the philosophical account, for instance, that there cannot be instantaneous action over a distance. The Einstein–Podolsky–Rosen experiment in physics shows an alleged case where perhaps there could be such action, under one interpretation. The philosopher then has an option of dropping the stricture from her metaphysics of causation. But is the interpretation a sound enough basis for dropping the stricture? There can be, thus, a toing and froing between the concrete and the abstract as our understanding develops. The goal is to find a theory that is satisfactory both scientifically and philosophically. This is what we mean by a reflective equilibrium.

In this book, we offer a sample of philosophically oriented discussions with which scientists should want to engage because they relate to the scientific discovery of causal connections. Should science construct causal theories, or only deal with data? How can our scientific studies have external validity, beyond the data set? And can we assume that the more data, the better is the evidence of causation? What do we

mean when we interpret a theory probabilistically? And why do we repeat experiments if they have already been shown to work once?

These questions show how normative philosophical considerations and empirical science can be intertwined and mutually informed. If treated by the philosopher in isolation, we will miss out on important evidence of how the world is observed to work. At the same time, science will be impoverished, lacking critical reflections over its conceptual, ontological, and methodological foundations. Thus, through engaging in metascience, there are good prospects of finding a complete and full understanding.

Uncovering the causal connections within nature is a key role for science. We must therefore consider how well the methods of science are suited for this task. Since all methods seem to reflect a set of deeper conceptual and ontological commitments concerning the nature of causation, we will have to make these commitments explicit and open for scrutiny. This is a philosophical task to which we now turn.

## 2

# Do We Need Causation in Science?

## 2.1 Causal Scepticism

There is a well-known, often quoted claim that causation has no place in science. We need to address this challenge as a matter of urgency for, if true, there would be no point in us continuing this book. There would be no causation in developed scientific theories for us to investigate. The claim is articulated by Bertrand Russell:

All philosophers, of every school, imagine that causation is one of the fundamental axioms or postulates of science, yet, oddly enough, in advanced sciences such as gravitational astronomy, the word 'cause' never occurs... 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. (Russell 1913: 193)

Russell's view is old but powerful and has retained support. There have been continuing interpretations of science—physics in particular—in which it is free of causation (Heisenberg 1959: 82, Feynman 1967: 147, Healey 1992: 193, Price and Corry 2007b, French 2014: 228).

How can this be so? Russell, and those who followed, argued that when you look at a successful and well-developed science such as physics, you find no claims of causal connection, such as regularities of same cause, same effect. What you find instead are equations that express co-variances of various magnitudes, e.g.  $F = ma$  or  $F = Gm_1m_2/d^2$ . Causation, as philosophers have traditionally conceived of it, has a number of features that we just do not see here. In particular, a causal connection is typically understood as directed and asymmetric. Causes produce their effects, rather than vice versa. Many say that causes must precede their effects in time. But we do not get this feature with equations, which exhibit symmetry instead. The equivalences they state can be read from left to right or right to left. Many sciences use differential equations: not just theoretical physics but also applied sciences such as engineering. Using these, we can calculate backwards as easily as forwards from any configuration. Hence, the future 'determines' the past just as much as the past 'determines' the future (Russell 1913: 201–2). In a science whose central tenets are articulated in the form of equations, therefore, there is no asymmetry represented. Allegedly, then, a physicist has no use of a concept of causation because it would commit us to asymmetric determinations that are just no longer believed by the experts to be features of the world.

What, then, should we say about the status of causation? In particular, how can we account for the common practice of thinking of the world in terms of causal connections? Even in an apparently scientific context we say, for instance, that heating an iron rod made it expand, the Fukushima nuclear power plant disaster was due to a tsunami caused by an earthquake at sea, or that an acceleration of a particular body occurred because of a collision. All of these look, whether explicitly or implicitly, like causal claims.

Those who follow the spirit of Russell's account usually admit that there is still plenty of causal talk, sometimes even by scientists, on informal occasions. But causation nevertheless does not appear in precise statements of scientific laws and should not, therefore, be deemed ultimately real. As causation is not a genuine part of the world, in serious science we should make no causal claims.

Causation, according to this view, could be seen as nothing more than 'folk science', which means 'a crude and poorly grounded imitation of more developed sciences' (Norton 2007: 12). Similarly, Ladyman et al. (2007: 3–4) think that causation is one of the pre-scientific notions that some have attempted to 'domesticate' into modern science, persisting with a common-sense metaphysics that has no basis in contemporary physics. Norton's argument is restricted not just to physics, though. He thinks that no developed science needs the notion of causation:

What I do deny is that the task of science is to find the particular expressions of some fundamental causal principle in the domain of each of the sciences. My argument will be that centuries of failed attempts to formulate a principle of causality, robustly true under the introduction of new scientific theories, have left the notion of causation so plastic that virtually any new science can be made to conform to it. Such a plastic notion fails to restrict possibility and is physically empty. This form of causal skepticism is ... motivated by taking the content of our mature scientific theories seriously. (Norton 2007: 12)

Norton contrasts causation with a respectable notion such as energy, which is to be found across a range of sciences. Among the things he says of it is 'there are innumerable processes that convert the energy of one science into the energy of another, affirming that it is all the same stuff. The term is not decorative; it is central to each theory' (Norton 2007: 14). The term causation, we can conclude, is mere decoration.

Can the causal realist respond to this sort of charge and answer the sceptic? We argue that there is a good answer and we should not abandon causation in science. Indeed, our conclusion will be even stronger than that. We argue that causation is vital for science: its bedrock, a pre-condition for its very existence. We proceed as follows. First we point out that the evidence from science far from warrants a firm conclusion that there is no causation. Second, we argue that science makes essential and ineliminable use of the notion of causation, with the centrality of intervention. And, third, empirical science is dependent on a notion of observation such that it would be rendered incoherent unless the reality of causation is accepted. Taken

together, these three points add up to a powerful argument in favour of there being causation in science.

## 2.2 Interpreting Science

First, suppose we restrict ourselves to physics, following the claims of Russell, Price (2007), and others in assuming this to be a science that invokes equations rather than causation. Perhaps this is the science where causation is most at risk; especially so if all other sciences ultimately reduce to physics (see Chapter 14). But here we should note that physics, as stated, is still open to philosophical interpretation, which is a matter of debate. Until such debates are resolved, are we really in a position to jettison so fundamental a concept as causation? Wouldn't we have to be very sure of our understanding of physics before we did so?

Physics is not yet complete and our understanding of it probably even less so. Are there really no asymmetries in reality? Some think the world exhibits a continuous increase in entropy—net energy dispersion—which provides an objective difference between past and future. Whether this is the case is controversial (Kutach 2002, Frisch 2007, Loewer 2007). Similarly, some think that physics allows for the possibility of backwards causation or 'retrocausality' (Price 1996, Price and Weslake 2009), thereby rendering the asymmetry of causation an illusion. Perhaps in some cases, effects can precede their causes, but this is also controversial. Norton (2007: 22) is at pains to point out that, even in Newtonian physics, some events are uncaused. Suppose we accept the claim. It still does not undermine what most causal realists believe. That causation is real does not require that every event has a cause. The realist's view is that causation is a part of reality in the cases where it does occur. All the above issues remain contested.

Another matter is how we are to understand the equations of physics. They look like a mathematical description or modelling of reality. But surely they should not be conflated with that reality. Although contemporary physics makes extensive use of maths, isn't that just a tool for representing certain structural features of the world? The world itself is not an equation, nor a number. Furthermore, an equation may deliberately represent only selected aspects of reality, for purposes of abstraction. A philosophical interpretation—outside of science itself—would then still be needed to supplement that representation in order to give us a complete understanding.

Examples might help. An equation fails to indicate asymmetry but that does not entail that there is none in the relationship between the things represented. Thus, although the formula  $F=ma$  is correct, this does not rule out some asymmetric causal relationships. We might causally intervene to increase the acceleration of an object by increasing the force on it, for instance. To take another case, there is an equation relating the length of a man's shadow on flat land, his height, and the position of the sun. From any two of these variables, we can calculate the third. Although this equation is in that respect symmetric, we also understand the direction