Studies in History and Philosophy of Biological and Biomedical Sciences xxx (2012) xxx–xxx

Contents lists available at SciVerse ScienceDirect

Studies in History and Philosophy of Biological and Biomedical Sciences

journal homepage: www.elsevier.com/locate/shpsc

Introduction

Phyllis Illari [a], Julian Reiss [b], Federica Russo [c,d]

[a] Department of Science and Technology Studies, University College London, Gower Street, London WC1E 6BT, United Kingdom
[b] Department of Philosophy, Durham University, 50 Old Elvet, Durham DH1 3HN, United Kingdom
[c] Center Leo Apostel, Vrije Universiteit Brussel, Krijgskundestraat 33, 1160 Brussels, Belgium
[d] Centre for Reasoning, School of European Culture & Languages, University of Kent, Canterbury, Kent CT2 7NX, United Kingdom

When citing this paper, please use the full journal title *Studies in History and Philosophy of Biological and Biomedical Sciences*

1. Causality in the sciences

The concept of cause is of extraordinary importance for the sciences. Scientists want to know the causes of phenomena because they want to be able to predict them, explain them, and gain control over them via interventions. In the second half of the twentieth century many researchers who were influenced by logical positivism, especially in the social sciences, tried to avoid the term 'cause' and its cognates. But much of their work always remained implicitly causal, and the concept experienced a philosophical revival towards the end of the century.

Today causality is one of the most fertile areas of research in the philosophy of science, as the papers in this special section attest. They demonstrate that the philosophy of causality goes well beyond the Humean questions, 'What is causality?' and 'How can we know?', and ranges across topics as varied as causal probabilities (Drouet), inferentialist semantics (Reiss), hierarchies among causal models (Hoover), the modal character of interventions (Reutlinger), the causal structure of mechanisms (Menzies), difference-making and mechanistic evidence for causal claims (Claveau), and statistical norms in causal attributions (Sytsma et al.).

These papers were first presented at the conference 'Causality in the Biomedical and Social Sciences', which was held in Rotterdam in 2010 and is part of the Causality in the Sciences (CitS) conference series.

2. The 'CitS' conference series

Causality in the Sciences or CitS (http://www.kent.ac.uk/secl/philosophy/jw/cits.htm) is a conference series that has been running since 2006, with the current steering committee of Isabelle Drouet (Paris), Phyllis Illari (UCL), Bert Leuridan (Ghent), Julian Reiss (Durham), Federica Russo (Brussels & Kent), Erik Weber (Ghent), and Jon Williamson (Kent). The series aims to promote fruitful dissemination of research by bringing together researchers actively working on causality in the UK, Europe, the US and elsewhere. Conferences regularly succeed in attracting researchers from scientific disciplines as well as from philosophy. To name but a few, computer scientists, researchers working in AI and bioinformatics, statisticians, and medical practitioners have regularly participated at recent meetings.

3. Why 'Causality in the Biomedical and Social Sciences (CiBaSS)'?

Conferences in the CitS series feature a different theme every year such as 'Causality and Probability' in 2006 and 2008, 'Causality and probability in the biomedical sciences' in 2006, and 'Mechanisms and Causality' in 2009. For the 2010 conference, which was organised by Julian Reiss, we decided to focus on 'Causality in the Biomedical and Social Sciences' because these sciences have been neglected in the causality literature until relatively recently, but there has been a flourishing of excellent work. This special section is broad and was kept broad on purpose, to allow us to discover what hot topics are emerging in the area, rather than impose a view of the direction of research. And indeed the finished set of papers shows the new diversity of really good work on biomedical and Social sciences.

4. The papers

In 'Causal reasoning, causal probabilities, and conceptions of causation' Isabelle Drouet examines a formal tool used to represent causal relationships, viz. causal probabilities. Causal probabilities are probabilities of causal effects: $P(c \triangleright e)$. Drouet takes up a paper by James Joyce in which Joyce argues that the values of causal probabilities do not depend on the metaphysics of causation, or more specifically on whether one conceives of causation in terms of interventions, of counterfactuals, or of probabilities. Drouet calls this the 'weak independence thesis' and distinguishes it from the

E-mail address: julian@jreiss.org (J. Reiss)

1369-8486/$ - see front matter © 2012 Published by Elsevier Ltd.
http://dx.doi.org/10.1016/j.shpsc.2012.09.001

following 'strong independence thesis': definitions of causal probabilities do not depend on whether one conceives of causation in terms of interventions, of counterfactuals, or of probabilities. Drouet continues to argue in favour of the strong independence thesis and to point out various difficulties with Joyce's argument in favour of the weaker thesis. Essentially, her argument is that causal probabilities under their various definitions cannot take the same values in general because they do not measure the same thing. This result has important ramifications for causal epistemology, which is used across the biomedical and social sciences, but also for causal decision theory.

In 'Causation in the sciences: An inferentialist account' Julian Reiss presents an alternative account of causation in the biomedical and social sciences according to which the meaning of causal claims is given by their inferential relations to other claims. Most accounts of causation that aim to provide semantics for causal claims (such as Hume's, Lewis's, Woodward's, Cartwright's etc.) are 'representationalist' in that they regard the meaning of causal claims as being given by what these claims represent – for instance, regularities, counterfactual dependencies, invariances under intervention or concrete causings. Inferentialism, by contrast, sees the meaning of claims as being constituted by their inferential relations with other claims. Reiss calls the set of claims that are thus inferentially related the 'inferential system' for a causal claim, and distinguishes between an 'inferential base' – the claims from which a causal claim is inferred – and the 'inferential target' – the claims that are inferred from a causal claim. Typical claims in the inferential base for a causal claim in the biomedical and social sciences are evidential claims such as claims about randomised trials, cohort studies, observational studies or expert judgement. Typical claims in its inferential target are claims about explanation, prediction, intervention and responsibility. Reiss then draws conclusions from this new way thinking about the nature of causal claims for the epistemology, semantics and metaphysics of causation.

In 'Causal structure and hierarchies of models' Kevin Hoover defends a perspectival realist account of causation in economics. Perspectival realism is the view according to which scientific descriptions capture only selected aspects of reality, and those aspects are not bits of the world seen as they are in themselves, but bits of the world seen from a distinctive human perspective. In Hoover's account the perspective is given by models that represent causal structure. Hoover here develops his earlier work on causal structure, itself a generalisation of Herbert Simon's approach to causation. The main contributions of this paper are a new notion of causal identity (specifically, two variables are said to be causally identical if and only if they share the same parameter set within a model; that is, when they are determined in a subsystem that cannot be separated into smaller subsystems) and the distinction of three types of modularity. The first type is related to the econometrician's notion of 'value free': a parameter is said to be value free when its range of values is unconstrained by the values other parameters take. The second type has received much attention from philosophers who work on interventionist conceptions of causality and consists in the independence of a causal relation between two variables from interventions that set other variables to particular values or break causal linkages. Hoover thinks that economic models are frequently not modular in this sense, but nonetheless are causally ordered – against what defenders of interventionist views of causation hold. The third type is most closely related to Hoover's perspectival realism. It takes an engineer's perspective: When a unit can be constructed out of parts, the parts can be considered modules at one level and the constructed unit a module at higher level. A module, in turn, is constructed by establishing conditions in which the conceptually distinct parts are causally identical. But if they are causally identical, they cannot be intervened on separately. This does not, however, prevent one from decomposing the complex system and analysing what appears as a module from one perspective as itself having internal structure from another perspective.

In 'Getting Rid of Interventions', Alexander Reutlinger examines the heart of the interventionist account of causality – interventions themselves. As Reutlinger says, Woodward's definitions of causal notions imply that the truth of certain causal statements requires the existence of at the least a possible intervention on the putative cause. Just to be clear, Woodward himself doesn't require interventions to be actual, or even physically possible, but rather something weaker than that. Reutlinger argues that Woodward's notion of an intervention is deeply problematic, by pushing a dichotomy. On the one hand, merely logically possible interventions are dispensable for giving an account of the meaning of causal statements. This is so because merely logically possible interventions fail to contribute non-trivially to the truth conditions of causal claims, being entirely restatable in terms of values of variables for the putative cause, so long as it is not a logical contradiction that the variables take the necessary values. On the other hand, merely logically possible interventions are indispensable, but the resulting account is inadequate. In brief, relying on merely logically possible interventions strongly conflicts with the standard approaches to the meaning of counterfactuals. If an existential claim about a merely logically possible intervention figures in the antecedent of an interventionist counterfactual, then this leads to the fatal result that this interventionist counterfactual is evaluated as false, although we would take these counterfactuals to be true. This is so because possible worlds as which logically possible interventions about some things happen, such as God intervening on the Big Bang, have very different laws of nature from the actual world, and such worlds are ruled out for evaluation of counterfactuals by the three standard approaches to the semantics of counterfactuals. Reutlinger notes that these are both compatible with the view that interventions are methodologically important, possibly indispensable, for discovering causal relations. Ultimately, he leaves it to the reader to decide which stance to adopt.

In 'The causal structure of mechanisms', Peter Menzies argues that hitherto the mechanisms literature has not given a full account of what is distinctive or essential to the causal structure of mechanisms, and so to mechanistic explanation. Menzies seeks to provide an account to remedy that by applying an interventionist account of causation within a structural equations framework. Menzies takes Craver's view as the one that pays most attention to causal structure, and builds on Craver's existing use of Woodward's work to argue that the interventionist account can be further used to fill holes in Craver's account, most importantly the formulation of the modularity constraint important to mechanistic explanation. Menzies holds that Craver's own account makes it clear that identifying the parts of mechanisms is important, usefully discussing the requisite experimental work, but does not tell us how the parts need to be causally organized to exhibit or constitute the phenomenon being explained. What is vital for this work, and so for building a mechanistic explanation, is to be able to affect the activity of a particular part without affecting those of other parts, except the effects. This, then, is the modularity constraint on mechanisms that Menzies identifies. Innovatively, he suggests this can be met by applying the term 'causal mechanism' to a system of causal relations described by a complete set of structural equations rather than a single structural equation. A single structural equation that holds invariantly under the appropriate range of actual and possible interventions describes a causal capacity of the system, rather than a mechanism. Ultimately, then, the causal structure of a mechanism is given a by a set of modular subcapacities whose sequential exercise has the input-output profile of the capacity to be explained. Menzies closes by further defending this

Please cite this article in press as: Illari, P., et al. Introduction. *Studies in History and Philosophy of Biological and Biomedical Sciences* (2012), http://dx.doi.org/10.1016/j.shpsc.2012.09.001

proposal by applying it to the mechanism of neurotransmitter release, and cognitive psychology approaches to reading abilities.

In 'The Russo–Williamson Theses in the social sciences: Causal inference drawing on two types of evidence', François Claveau examines two theses formulated by Federica Russo and Jon Williamson in their study of causal inference in the health sciences. Claveau asks whether the theses generalise to the social sciences. Specifically, he assesses their validity against evidence from a specific case the economics, viz. research on the institutional determinants of the aggregate unemployment rate. What he calls the 'First Russo–Williamson Thesis' is that a causal claim can only be established when it is jointly supported by difference-making and mechanistic evidence. Claveau argues that this thesis does not hold for the economics case study. While unemployment researchers generally draw extensively on both difference-making and mechanistic evidence, one specific causal claim Claveau examines is regarded as having been established by the community of researchers, even though it is exclusively supported by mechanistic evidence. The 'Second Russo–Williamson Thesis' is that standard – monist – accounts of causality (such as regularity, probabilistic or process accounts) fail to support the dual epistemology highlighted in the first thesis. Against the second thesis, Claveau argues that a counterfactual-manipulationist account of causality can make sense of the typical epistemic strategy employed by social researchers. He shows that it is in fact just an instance of the common more general strategy of seeking to support scientific claims with a broad variety of evidence.

In 'Two Types of Typicality: Rethinking the Role of Statistical Typicality in Ordinary Causal Attributions', Justin Sytsma, Jonathan Livengood, and David Rose examine the sensitivity of causal attributions to statistical norms. The psychology literature agrees that causal attributions are sensitive to both statistical and prescriptive norms, so, for example, people's judgments of Jack's causal responsibility for the secretaries running out of pens on their desk depends both on whether it is usual or unusual to take their pens (statistical norm) and on whether there are any rules against taking their pens (prescriptive norm). However, Sytsma, Livengood and Rose distinguish two kinds of statistical norms: population-level statistical norms, which are about what is usual in a population of people, and agent-level statistical norms, which are about what is usual for a single person. Sytsma, Livengood and Rose then argue for what they call the responsibility view, whereby statistical norms do not impact causal attributions directly, but by reflecting people's moral judgments. They make two predictions about causal attributions on its basis: ordinary causal attributions for the Pen Case are insensitive to population-level statistical norms, instead they are sensitive to agent-level typicality, not atypicality. Specifically, in cases like the Pen Case people will be more likely to say that an agent who behaves in a way that is in the agent-level statistical norm caused a bad outcome than to say that an agent who behaves in a way that is out of the agent-level statistical norm caused the same bad outcome. They then describe their experimental work, and its results, which support these predictions, undermining the received consensus.

Ultimately, we hope that the diverse work here means many people will find something of interest. Enjoy the ride!

Acknowledgements

We would like to thank the following for acting as referees for the papers of this special section: William Bechtel, James Bogen, Alex Broadbent, Nancy Cartwright, Dominique Chu, Carl Craver, Damien Fennell, José Garcia Encinas, Donald Gillies, Stuart Glennan, Till Grüne-Yanoff, Christopher Hitchcock, James Joyce, Joshua Knobe, Mainard Kuhlmann, Bert Leuridan, Peter Menzies, Alessio Moneta, Robert Northcott, Iñaki San Pedro, James Tabery, Dingmar van Eck, Erik Weber and Jon Williamson. Julian Reiss would like to thank the Faculty of Philosophy of Erasmus University Rotterdam and the Erasmus Trustfonds for financial support of the CiBaSS conference.

Please cite this article in press as: Illari, P., et al. Introduction. *Studies in History and Philosophy of Biological and Biomedical Sciences* (2012), http://dx.doi.org/10.1016/j.shpsc.2012.09.001