Causality in the Sciences
Edited by Russo, Williamson and Illari
Oxford University Press, 2011, pp. 952, £95
ISBN: 978-0-19-957413-1
doi:10.1017/S0031819113000193

For many philosophers working outside the field, causality is an issue for metaphysics. Indeed, this is how causality is often investigated. Yet a recent trend challenges this tradition, asking epistemological and methodological questions, as well. At the same time, causal language has been witnessing resurgence in many sciences. Scientists formerly comfortably only with probabilistic language are increasingly comfortable speaking of 'causes'. *Causality in the Sciences* brings together papers from these new approaches to causation. It includes sciences often left out of the discussion and considers scientific attempts to grapple with the notion of cause. It is very much unlike any other collection on causation and thus a worthwhile contribution.

The volume is mammoth. It contains 42 original articles from 71 contributors, complete with introduction and index. A handful of invited pieces complement the bulk of papers, which were submitted to an open call. As a result, contributing authors range from (academic) household names, like Judea Pearl and Nancy Cartwright, to relative newcomers. The list of contributors is in itself a valuable item, for it reveals a spectrum of philosophers, psychologists, economists, and all manner of other scientists interested in causality.

The editors attribute the new field of causation research to the rehabilitation of causation through its mathematisation. Of particular relevance are the developments of Bayesian Net methods for causal reasoning and structural equation models for causal relations. The importance of these tools comes to the fore in the penultimate section, which focuses almost exclusively on these formal methods. Yet there is more to the growth of this new field. I would add that the popularity of the 'new mechanist' approach to causation has helped greatly. A mechanist view of nature was prominent in the 17^{th} century, according to which phenomena were viewed as the product of mechanical interactions of matter. Though this view fell out of fashion, a 'new mechanistic philosophy' has recently emerged. The new mechanists are inspired by the observation that science often proceeds by investigating and modelling mechanisms.

https://doi.org/10.1017/S0031819113000193 Published online by Cambridge University Press

William Wimsatt[1] characterized biology thusly, while Bechtel and Richardson[2] described mechanistic explanation across the sciences. Machamer, Darden, and Craver built on this and other work, seeking to clarify the notion of mechanism and situate it in a theory of causation.[3] The number of papers in this volume touching on the concept of mechanism, slightly more 'pro' than 'contra', is testament to the impact of this perspective. The editors are certainly no strangers to this approach; though the title of the volume is *Causality in the Sciences*, it is the product of two research projects each explicitly focussed on mechanisms: 'Mechanisms and causality' and 'Causality across the levels: Biomedical mechanisms and public health policies'. It is unsurprising that the new mechanist approach forms a backdrop for much of the volume. Mechanisms take centre stage in the final section of the volume, where a discussion of the metaphysics of mechanisms begins to emerge. This is a topic hitherto neglected, with which many will be pleased to engage.

The volume begins with a comprehensive introductory 'manifesto', in which the editors trace the history of the new approach to causation, and introduce each section. After the introduction come five sections on causality in individual branches of science: Health Sciences; Psychology; Social Sciences; and Computer Science, Probability, and Statistics. These are capped with the discussion of Causality and Mechanisms.

Of the individual sections, those on Health Sciences and Computer Science, Probability and Statistics are the most unified, easily standing on their own as coherent contributions to their fields. As for the other sections, the best way to appreciate them is to look for intersections across and between sciences. Individual readers will be attracted to different exchanges, but of particular interest to me were a pair of papers by Ken Aizawa & Carl Gillet and Tudor Baetu, examining multiple realisation from the standpoint of neuropsychology and molecular biology.

As a whole, the section on Health Sciences suggests that the relationship between data and theory needs re-examination. R. Paul

[1] Wimsatt, W. (1972). 'Complexity and Organization', in Kenneth F. Schaffner and Robert S. Cohen (eds), *PSA 1972, Proceedings of the Philosophy of Science Association* (Dordrecht: Reidel), 67–86.

[2] Bechtel, W. & Richardson, R (1993). *Discovering Complexity: Decomposition and Localization as Strategies in Scientific Research*, Princeton, NJ: Princeton University Press.

[3] Machamer, P., Dardeny, L. & Craver, C. (2000). 'Thinking about mechanisms', *Philosophy of Science* **67**, 1–25.

https://doi.org/10.1017/S0031819113000193 Published online by Cambridge University Press

Thompson sets the pace, arguing that Randomized Control Trials (RCTs), the 'gold standard' of evidence-based medicine, are on their own insufficient for causal claims. They can at best test predictions of a theory. Similar issues are dealt with in articles by Alex Broadbent, Harold Kinkaid, and Donald Gillies. Burt Leuridan and Erik Weber round off the section nicely, setting the very pragmatic task of evaluating the International Agency for Research in Cancer (IARC) guidelines for assessing carcinogenicity. This exercise in applied philosophy of science flexes its muscles by exposing some limitations of the mechanist programme.

These contributions could be gainfully set-against Nancy Cartwright's paper, which seems slightly incongruous later in the volume. Much of Cartwright's recent work addresses methodology and the usefulness of causal knowledge.[4] Consider the transition from experimental set-up to real world. The standard approach is to ensure that experiment resembles the real world in as many ways as possible. This is notoriously difficult, and much has been written about how to ensure this ideal is met. According to Cartwright, we should seek an altogether better framework for attaining causal knowledge, which focuses more on use. She advocates the studying of causal capacities, or what Mill called 'tendencies'.[5] Where traditional causal claims are empirical facts about specific situations, capacities are metaphysical facts about the objects themselves. Think of them as irreducible causal powers. Capacities may not be efficacious in certain situations, but as facts about objects they are more useful as causal knowledge than facts about contexts. In the health sciences, Cartwright's proposal might inform an alternative to RCTs, which satisfies Thompson's concerns.

Though Cartwright's piece serves as an excellent counter-point to many of the articles in this volume and provides an introduction to thinking about capacities, the reader is left wanting more. As elsewhere, Cartwright makes a compelling case for a focus on capacities, but leaves us without a formal method for studying them. The sceptical reader may wonder whether Cartwright's proposal is in fact practicable, and if so, how. That said, Cartwright has here and elsewhere acknowledged this methodological shortfall. Perhaps an ambitious scholar will soon take up the task.

[4] Cartwright, N. *Hunting Causes and Using Them*. (Cambridge UK: Cambridge University Press, 2007).

[5] Cartwright, N. *Nature's Capacities and Their Measurement*, (Oxford: Clarendon Press, 1989).

https://doi.org/10.1017/S0031819113000193 Published online by Cambridge University Press

The section on COMPUTER SCIENCE, PROBABILITY, AND STATISTICS is very much the stylistic outlier in the volume, yet the issues addressed are central to the project. The tasks for the formal study of causation are now threefold: perfect formal methods, strengthen requisite concepts, and explore philosophical implications. Papers concerning the automated learning of causal models, the interpretation of statistical significance, and the notion of causal strength provide an insightful mix of methodological and conceptual work. For many, however, the highlight of the section will be the contribution from Judea Pearl, which combines methodological and conceptual considerations in the service of unpacking the philosophical implications of his earlier work. Pearl's probabilistic modelling of causation[6] and his Structural Causal Model (SCM) approach[7] have had an immense impact on the field. It is no surprise that his work appears in every article in this section. Here, Pearl unpacks SCM to show how this formal structure informs a comprehensive theory of causation. Specifically, he shows how this theory can provide information about counterfactual conditionals and indirect effects. The uninitiated reader should not be put-off by the formalisms in Pearl's work; Pearl here takes a noticeably didactic approach, helping the reader along the way. This piece is worth the effort, as it will help the reader understand how general concerns about causation are translated into and resolved by formal systems like SCM. Of particular help is the patient way in which Pearl introduces and highlights the utility of schematic representations of causal systems. This section, with Pearl's contribution in particular, should open the sometimes-daunting formal study of causation to a new audience.

In their introduction, the editors flag up a convergence of topic between papers by Ken Aizawa & Carl Gillet and Tudor Baetu. Though each feels very out-of-place in its respective section, PSYCHOLOGY and NATURAL SCIENCES, read together, they provide an excellent discussion of multiple realisation and its implications for explanatory reductionism.

Those familiar with neuroscience will not be surprised by Aizawa and Gillet's claim that the objects of psychology are often not borne out by neuroscientific findings. What makes their analysis of interest is that they use the case of colour perception to illustrate how scientific practice may not reflect the standard philosophical treatment

6 Pearl, J. *Probabilistic Reasoning in Intelligent Systems* (San Mateo, CA: Morgan Kaufmann, 1988).

7 Pearl, J. *Causality: Models, Reasoning, and Inference* (New York: Cambridge University Press, 2000).

https://doi.org/10.1017/S0031819113000193 Published online by Cambridge University Press

of multiple-realization. The incumbent view holds that, in cases where a single higher-level category is found to be multiply realised, the explanatory strategy is to split the higher category corresponding to different lower-level realisers. This would mean splitting the higher-level category of 'normal colour perception' into many different types, corresponding to the many variations in lower-level properties (eg. photo pigments) which give rise to slightly different perception. Yet this has not occurred; the approach actually varies. Scientists either recognize variation in a single higher-level category or split the category in two. Aizawa and Gillet credit this decision to the recognition of 'weak' and 'strong' realisers, respectively. But Aizawa and Gillet leave us wondering whether this strategy is a good one. It may seem that the distinction between weak and strong realisers is metaphysically arbitrary, corresponding solely to the aims of the scientists.

A partial answer to this concern may come from Baetu's treatment of the alleged explanatory reduction of classical to molecular genetics. Baetu considers the famous case of Huntington's disease, where classical approaches could not explain why 5% of inherited cases fail to show the standard pathology. As it happens, Huntington's results from mutant repeats in a crucial gene, where the number of repeats corresponds to degrees of neuronal decay. The 5% are fortunate to have a small number of mutations. This has led to the recognition of gradation in the disease. Contrast this with the cases of Marfan, Loeys-Dietz, and Ehlers-Danlos syndromes. At the higher-level, these three autoimmune diseases are difficult to distinguish, often being classified as one. Yet it turns out that the three involve different mutations impacting one of two mechanisms, leading scientists to separate the disease category into three unique ones. Though not Baetu's concern, there is an important difference between these two cases. The Huntington's case evidences weak realisers, while the autoimmune cases are strong realisers. Yet what makes them so is not just that scientists decided to split the autoimmune categories but maintain the Huntington's one; rather, the weak realiser reflects gradation in the mechanism responsible for the disease, whereas the strong realiser reflects genuinely different mechansisms. Baetu's discussion provides not only an additional disciplinary perspective, but also details that help render Aizawa and Gillet's work metaphysically respectable. Of course, one might still wonder whether distinctions between mechanisms are natural or nominal, but for this the reader can take recourse to the final section of the volume.

The 'new mechanists' have tended to avoid discussion of metaphysics, favouring more practical matters. Where the mechanists of the

https://doi.org/10.1017/S0031819113000193 Published online by Cambridge University Press

17th century advanced a thesis about the structure of reality, the new mechanists focus on the epistemic utility of mechanisms, occasionally drawing inferences about the world from epistemic success. As Williamson and Illari claim, however, mechanisms *must* be real in order to underwrite explanations. They believe this unstated assumption to be present in most mechanist literature. Stathis Psillos tackles these implicit metaphysical theses head-on. Psillos' tactic is to situate the new focus on mechanisms in relation to the old. He introduces the 17th century conceptions along with critiques by Poincaré and Hegel, both of whom agreed with a mechanist world-view but believed the selection of one mechanistic description over another was metaphysically arbitrary. They claim that a finite set of parts can be arranged into many more mechanisms than we might like, and there is nothing about the mechanisms themselves which privileges one description over another. Scientists and philosophers so privilege by appealing to extrinsic properties: it is only when we have an end in mind that one mechanism becomes the correct description. We must therefore concede that the unification of a given group of parts into a coherent mechanism is a task performed by humans, rather than nature. Williamson and Illari acknowledge that mechanisms are functionally individuated, but seem not to appreciate the scope of Psillos' claims. Consider their discussion:

> 'The same object individuated structurally, such as the heart, can have different functions according to the description of the system it is in. It might have the function of pumping blood when considered as part of the circulatory system, or the function of making a thump-thump noise when considered as part of a system for comforting a newborn baby.' (826)

However, Psillos (via Hegel) raised a much stronger point. 'The heart' is not a metaphysically privileged mechanism; it is only 'the heart' when we situate it in a specific explanatory framework. The mechanism for 'thump-thump' would include a different sub-set of parts than the mechanism for pumping blood. The former involves fewer facts about the exchange of oxygen and more facts about the acoustic properties of the sternum. There *is* no physical description of the heart *until* you specify an end. The mechanism of the heart does not make it a real object; it is one possible description privileged only by human interest in explaining specific phenomena. This should trouble the mechanist. After all the, mechanisms require objects for their very constitution!

The mechanist project is intriguing and fruitful. This exchange shows that proponents are prepared to move beyond traditional

https://doi.org/10.1017/S0031819113000193 Published online by Cambridge University Press

concerns and begin thinking about the metaphysics of mechanisms, a topic sure to draw great interest. Future work might explore whether anti-realist considerations affect the epistemic utility of mechanist explanation.

The re-emergence of 'cause' talk in the sciences generates a series of interrelated tasks. The first is descriptive, determining what roles causation plays in scientific practices, to what end, and how; the second task is evaluative, appraising these activities and their ends; and the third is revisionary. In this volume, the reader will find articles dealing with each of these questions in all manner of sciences. Some articles in this volume struggle with this division of labour, mixing descriptive accounts with normative claims; but, where authors succeed, the reader comes away greatly enriched. *Causality in the Sciences* is a noteworthy achievement in a successful and highly productive research programme. Williamson, Russo, and Illari, among others, have worked tirelessly to shift the boundaries and methodologies of traditional philosophical treatments of causation, seeking a new take on an old problem. The breadth of articles and range of contributors in this weighty volume are testament to their success.

Jordan Bartol
phjnb@leeds.ac.uk
This review first published online 18 March 2013

Personhood, Ethics and Animal Cognition: Situating Animals in Hare's Two-Level Utilitarianism
By Gary E. Varner
Oxford and New York: Oxford University Press, 2012, pp. xiv + 317
ISBN: 978-0199758784
doi:10.1017/S0031819113000363

Gary Varner, who used to be, like the current reviewers, an ethical biocentrist, now defends Harean prescriptivism, two-level utilitarianism, and sentientism, and in this book applies these stances to animal ethics, as well as to ethical principles in general. As his Introduction discloses, Varner feels impressed by large areas of Richard Hare's thinking, not least because much of it inspired that of Peter Singer. In this work Varner seeks to supplement the work of Hare and of Singer by discussing ethical principles appropriate to the treatment of animals, embodying his distinction between persons, near-persons and sentient non-persons.

https://doi.org/10.1017/S0031819113000193 Published online by Cambridge University Press