Review of *Physical Causation*, by Phil Dowe (Cambridge: Cambridge University Press, 2000). ISBN 0-521-78049-7

Reviewer: Katherine Hawley, Department of Logic and Metaphysics, University of St Andrews, St Andrews, Fife, KY16 9AL, Scotland.
Email: kjh5@st-andrews.ac.uk Forthcoming in *MetaScience*.

Phil Dowe begins *Physical Causation* by outlining his metaphilosophical stance, distinguishing between the task of analysing our concept of causation, and the task of offering an "empirical analysis" of "what causation in fact *is* in the actual world" (p.3). Dowe's energies are to be devoted to the latter task, although he regards the former as also legitimate. Thus he is not restricted by what we mean by 'causation', and nor is he obliged to offer more than a contingent account: causation may be different things in different worlds. It is possible for philosophers to perform empirical analysis, so long as their work is informed by scientific results. Dowe does indeed draw upon "hints taken from science" (p.12), although the level of detail here neither greatly constrains his theorising nor makes the book inaccessible to scientifically-hazy metaphysicians.

In chapter two we are taken at a canter through some broadly Humean accounts of the nature of causation. Each has its flaws: a straightforward constant-conjunction account cannot accommodate indeterministic causation (where the occurrence of an event of the first kind is not sufficient for the occurrence of an event of the second); Lewis-style counterfactual accounts are vulnerable to the converse problem of overdetermination (where the cause is not necessary for the effect); any identification of causation with probability-raising falls to the fact that a cause may lower the probability of its effect (by bringing it about in some unlikely way, whilst preventing it from occurring in some safer way). Proponents of these various accounts do not surrender to such attacks, and indeed Dowe discusses various possible responses, before rejecting them all. Nevertheless, the relative clarity and brevity of Dowe's treatment provides us with an elegant, systematic overview of some rather dense literature.

'Transference' theories—according to which a cause transfers energy, or perhaps momentum, to its effect—are an improvement over Humean theories, argues Dowe, but they fall short in three respects. First, they apparently cannot encompass 'immanent' causation, in which an object causes its own continuing existence (or continuing motion) in the absence of any *transfer* of energy. Second, they require an unfeasibly robust notion of identity over time for individual quantities of energy (or of momentum). Third, such accounts cannot distinguish between giving and receiving a quantity except by plugging the direction of time into the direction of causation—the causing object is the one that has the quantity at the earlier time. Dowe wants to allow for effects that precede their causes, as we shall see.

Russell, and, later, Wesley Salmon, base accounts of causation upon the notion of a 'causal process', where for Salmon a causal process is one which can transmit a mark. A key question here is whether 'transmit' and 'mark' can be spelt out in ways that do not presuppose an account of causation. Dowe is sympathetic, yet not entirely satisfied. In short, he is attracted to the idea of taking processes as basic, whilst he rejects Salmon's 'mark transmission' characterisation of the difference between causal and non-causal processes. He also finds something appealing about the transferred quantity accounts, although the requirement of *transfer* seems too restrictive. Out of these considerations emerges Dowe's Conserved Quantity account of causation.

Here goes. Spacetime regions can be divided into those that contain objects—particles, chairs, shadows, water-waves—and those that are merely arbitrary or gerrymandered. Regions that are worldlines of objects are processes. Amongst processes, causal processes are those in which the object possesses a conserved quantity, although the amount of the quantity possessed by the object may vary as it interacts with other objects. (A conserved quantity is a quantity, like mass–energy or charge, that is governed by a conservation law.) So the worldline of a chair is a causal process, whilst the worldline of a shadow is a process, but not a causal process. Finally, a causal interaction is an intersection of worldlines that involves exchange of a conserved quantity.

Unlike transference theories, the Conserved Quantity theory can account for immanent causation. And Dowe does not need to track the identity of quantities across time—in the exchange of a conserved quantity, the amount of energy (for example) possessed by one object goes down, whilst the amount possessed by the other goes up, but nothing turns on whether the very same piece of energy has been lost by the former and gained by the latter. Like Salmon, Dowe focuses on causal processes, but he replaces the obscurer notions of 'mark' and 'transmit' with the brighter notion of 'conserved quantity'. As Dowe admits, he relies upon our grasping a distinction between genuine objects and mere gerrymandered worldlines. On pain of circularity, he cannot use the common strategy of picking out genuine objects as those which are sustained by internal causal connections. Indeed, he wants to recognise objects, like light-spots and shadows, that are not sustained by such internal causal connections. Nevertheless, he thinks that if we are to take a notion as basic, then that of identity through time is not such a bad choice.

Dowe's theory is clearly and carefully advocated, connections to rivals are noted, and traps are avoided. I was unsettled by an explanatory gap, though Dowe, presumably, will reply that his 'empirical analysis' is not bound to provide much in the way of explanation. The unanswered question is this: *why* are all and only conserved quantities involved in causation? Recall the picture: causal objects like hawks and handsaws are distinguished from 'noncausal' objects like shadows and light-spots by the fact that the former, and only the former, possess conserved quantities. Noncausal objects may have velocity and shape, for example, but not energy, momentum or charge. Why is this? Is there a connection between the fact that an object possesses quantities which cannot be created and destroyed, and the fact that such an object can make a difference to other objects, and to its own future existence? Is the necessity of conservation somehow the ground of the necessity of causation, and, if so, is this because conservation laws are in some way more basic than other laws?

Dowe is explicitly neutral (pp. 95–6) about what causation would be in a possible world in which, say, mass-energy was not conserved by some novel, alien quantity was governed by a conservation law. (This scenario prompts interesting thoughts about the possibility of irreducibly mental causation.) And, in so doing, he certainly stays within the job description of the empirical analyst. But there is plenty of scope here for others to speculate.

Chapters six and seven elaborate the Conserved Quantity theory so that Dowe can deal with causation by prevention and omission, and so that he can identify causes and effects (not just any two events connected by causal processes are related as cause and effect). The work on prevention and omission, in particular, should be of interest even to those who reject Dowe's overall picture. In brief, negative 'causation', by prevention or omission, is not genuine causation, but for practical purposes we are justified in treating it as causation, since it is a matter of real (positive) causation plus some counterfactual truths about positive causation.

In his final chapter, Dowe presupposes that Bell-type phenomena in quantum mechanics reveal the existence of backwards-in-time causation (I didn't find the macroscopic model of the phenomena especially illuminating, but others may differ), then he uses the phenomena as a test case for a variety of philosophical accounts of causal asymmetry. Dowe opts for a version of the 'conjunctive fork' account, in which (very roughly) we distinguish effects from causes by statistically characterising the fact that when a cause has joint effects, the two effects are neither independent of each other, nor directly dependent upon each other. Joint effects of a common cause are differently related than joint causes of a common effect.

This means that, if Bell-type situations exemplify backwards causation, then, as the later measurement causes the earlier state of the system, there will be some other early event that is suitably correlated with the early state of the system, and which is caused by the later measurement. Thus, says Dowe, by searching for such an event, we can empirically test the hypothesis that there is backwards-in-time causation. This intriguing suggestion is an apt place at which to end this valuable, thought-provoking investigation into the empirical nature of causation.