Introduction

Look at what economists are saying. 'Changes in the real GDP unidirectionally and significantly Granger cause changes in inequality.'[1] Alternatively, 'the evolution of growth and inequality must surely be the outcome of similar processes' and 'the policy maker . . . needs to balance the impact of policies on both growth and distribution'.[2] Until a few years ago claims like this – real causal claims – were in disrepute in philosophy and economics alike and sometimes in the other social sciences as well. Nowadays causality is back, and with a vengeance. That growth causes inequality is just one from a sea of causal claims coming from economics and the other social sciences; and methodologists and philosophers are suddenly in intense dispute about what these kinds of claims can mean and how to test them. This collection is for philosophers, economists and social scientists or for anyone who wants to understand what causality is, how to find out about it and what it is good for.

If causal claims are to play a central role in social science and in policy – as they should – we need to answer three related questions about them:

What do they mean?
How do we confirm them?
What use can we make of them?

The starting point for the chapters in this collection[3] is that these three questions must go together. For a long time we have tended to leave the first to the philosopher, the second to the methodologist and the last to the policy consultant. That, I urge, is a mistake. Metaphysics, methods and use must march hand in hand. Methods for discovering causes must be legitimated by showing that they are good ways for finding just the kinds of things that causes are; so too the conclusions we want to draw from our causal claims, say for planning and policy, must be conclusions that are warranted given our account of what causes are. Conversely, any account of what causes are that does not dovetail with what we take to be our best methods for finding them or the standard

---

[1] Assane and Grammy (2003), p. 873. [2] Lundberg and Squire (2003), p. 326.
[3] Many of these chapters have previously been published; about one-third are new. For original places of publication, see the acknowledgements.

1

uses to which we put our causal claims should be viewed with suspicion. Most importantly –

Our philosophical treatment of causation must make clear why the methods we use for testing causal claims provide good warrant for the uses to which we put those claims.

I begin this book with a defence of causal pluralism, a project that I began in *Nature's Capacities and their Measurement*,[4] which distinguishes three distinct levels of causal notions, and continued in the discussions of causal diversity in *The Dappled World*.[5] Philosophers and economists alike debate what causation is and, correlatively, how to find out about it. Consider the recent *Journal of Econometrics* volume on the causal story behind the widely observed correlations between bad health and low status. The authors of the lead article,[6] Adams, Hurd, McFadden, Merrill and Ribeiro, test the hypothesis that socio-economic status causes health by a combination of the two methods I discuss in part II: Granger causality, which is the economists' version of the probabilistic theory of causality that gives rise to Bayes-nets methods, and an invariance test. Of the ten papers in the volume commenting on the Adams et al. work, only one discusses the implementation of the tests. The other nine quarrel with the tests themselves, each offering its own approach to how to characterize causality and how to test for it.

I argue that this debate is misdirected. For the most part the approaches on offer in both philosophy and economics are not alternative, incompatible views about causation; they are rather views that fit different kinds of causal systems. So the question about the choice of method for the Adams et al. paper is not 'What is the "right" characterization of causality?' but rather, 'What kind of a causal system is generating the AHEAD (Asset and Health Dynamics of the Oldest Old) panel data that they study?'

Causation, I argue, is a highly varied thing. What causes should be expected to do and how they do it – really, what causes are – can vary from one kind of system of causal relations to another and from case to case. Correlatively, so too will the methods for finding them. Some systems of causal relations can be regimented to fit, more or less well, some standard pattern or other (for example, the two I discuss in part II) – perhaps we build them to that pattern or we are lucky that nature has done so for us. Then we can use the corresponding method from our tool kit for causal testing. Maybe some systems are idiosyncratic. They do not fit any of our standard patterns and we need system-specific methods to learn about them. The important thing is that there is no single interesting characterizing feature of causation; hence no off-the-shelf or one-size-fits-all method for finding out about it, no 'gold standard' for judging causal relations.[7]

---

[4] Cartwright (1989).   [5] Cartwright (1999).   [6] Adams et al. (2003).
[7] See John Worrall (2002) on why randomized clinical trials are not a gold standard.

Part II illustrates this with two different (though related) kinds of causal system, matching two different philosophical accounts of what causation is, two different methodologies for testing causal claims and two different sets of conclusions that can be drawn once causal claims are accepted.

The first are systems of causal relations that can be represented by causal graphs plus an accompanying probability measure over the variables in the graph. The underlying metaphysics is the probabilistic theory of causality, as first developed by Patrick Suppes. The methods are Bayes-nets methods. Uses are licensed by a well-known theorem about what happens under 'intervention' (which clearly needs to be carefully defined) plus the huge study of the counterfactual effects of interventions by Judea Pearl. I take up the question of how useful these counterfactuals really are in part III.

In part II, I ask 'What is wrong with Bayes nets?' My answer is really, 'nothing'. We can prove that Bayes-nets methods are good for finding out about systems of causal relations that satisfy the associated metaphysical assumptions. The mistake is to suppose that they will be good for all kinds of systems. Ironically, I argue, although these methods have their metaphysical roots in the probabilistic theory of causality, they cannot be relied on when causes act probabilistically. Bayes-nets causes must act deterministically; all the probabilities come from our ignorance. There are other important restrictions on the scope of these methods as well, arising from the metaphysical basis for them. I focus on this one because it is the least widely acknowledged.

The second kind of system illustrated in part II is systems of causal relations that can be represented by sets of simultaneous linear equations satisfying specific constraints. The concomitant tests are invariance tests. If an equation represents the causal relations correctly, it should continue to obtain (be invariant) under certain kinds of intervention. This is a doctrine championed in various forms by both philosophers and economists. On the philosophical side the principal advocates are probably James Woodward and Daniel Hausman; for economics, see the paper on health and status mentioned above or econometrician David Hendry, who argues that causes must be *super*exogenous – they must satisfy certain probabilistic conditions (exogeneity conditions) and they must continue to do so under the policy interventions envisaged. (I discuss Hendry's views further in chs. 4 and 16.)

My discussion in part II both commends and criticizes these invariance methods. In praise I lay out a series of axioms that makes their metaphysical basis explicit. The most important is the assumption of the priority of causal relations, that causal relations are the 'ontological basis' for all functionally true relations, plus some standard assumptions (like irreflexivity) about causal order. 'Two theorems on invariance and causality' first identifies a reasonable sense of 'intervention' and a reasonable definition of what it means for an equation to 'represent the causal relations correctly' and then proves that the methods are

matched to the metaphysics. Some of the uses supported by this kind of causal metaphysics are described in part I.

As with Bayes nets, my criticisms of invariance methods come when they overstep their bounds. One kind of invariance at stake in this discussion sometimes goes under the heading 'modularity': causal relations are 'modular' – each one can be changed without affecting the others. Part II argues that modularity can – and generally does – fail.

I focus on these two cases because they provide a model of the kind of work I urge that we should be doing in studying causation. Why is it that I can criticize invariance or Bayes-nets methods for overstepping their bounds? Because we know what those bounds are. The metaphysical theories tell us what kinds of system of causal relations the methods suit, and both sides – the methods and the metaphysics – are laid out explicitly enough for us to show that this is the case. The same too with the theorems on use. This means that we know (at least 'in principle') when we can use which methods and when we can draw which conclusions.

Part III of this book looks at a number of economic treatments of causality. The chapter on models and Galilean experiments simultaneously tackles causal inference and another well-known issue in economic methodology, 'the unrealism of assumptions' in economic models. Economic models notoriously make assumptions that are highly unrealistic, often 'heroic', compared to the economic situations that they are supposed to treat. I argue that this need not be a problem; indeed it is necessary for one of the principal ways that we use models to learn about causes.

Many models are thought experiments designed to find out what John Stuart Mill called the 'tendency' of a causal factor – what it contributes to an outcome, not what outcomes will actually occur in the complex world where many causes act together. For this we need exceptional circumstances, ones where there is nothing else to interfere with the operation of the cause in producing its effect, just as with the kinds of real experiment that Galileo performed to find out the effects of gravity. My discussion though takes away with one hand what it gives with the other. For not all the unrealistic assumptions will be of this kind. In the end, then, the results of the models may be heavily overconstrained, leading us to expect a far narrower range of outcomes than those the cause actually tends to produce.

The economic studies discussed in part III themselves illustrate the kind of disjointedness that I argue we need to overcome in our treatment of causality. Some provide their own accounts of what causation is (economist/methodologist Kevin Hoover and economists Steven LeRoy and David Hendry); others, how we find out about it (Herbert Simon as I reconstruct him and my own account of models as Galilean experiments); others still,

what we can do with it (James Heckman and Steven LeRoy on counterfactuals). The dissociation can even come in the interpretation of the same text. Kevin Hoover (see ch. 14, 'The merger of cause and strategy: Hoover on Simon on causation') presents his account as a generalization to non-linear systems of Herbert Simon's characterization of causal order in linear systems. My 'How to get causes from probabilities: Cartwright on Simon on causation' (ch. 13) provides a different story of what Simon might have been doing. The chief difference is that I focus on how we confirm causal claims, Hoover on what use they are to us.

The turn to economics is very welcome from my point of view because of the focus on use. In the triad metaphysics, methods and use, use is the poor sister in philosophic accounts of causality. Not so in economics, where policy is the point. This is why David Hendry will not allow us to call a relation 'causal' if it slips away in our fingers when we try to harness it for policy. And Hoover's underlying metaphysics is entirely based on the demand that we must be able to use causes to bring about effects.

Perhaps it seems an unfair criticism of our philosophic accounts to say they are thin on use. After all one of our central philosophic theories equates causality with counterfactuals and another equates causes with whatever we can manipulate to produce or change the effect. Surely both of these provide immediate conclusions that help us figure out which policies and techniques will work and which not? I think not. The problem is one we can see by comparing Hoover's approach to Simon with mine. What we need is to join the two approaches in one, so that we simultaneously know how to establish a causal claim and what use we can make of that claim once it is established.

Take counterfactuals first. The initial David Lewis style theory[8] takes causal claims to be tantamount to counterfactuals: C causes E just in case if C had not occurred, E would not have occurred. Recent work looks at a variety of different causal concepts – like 'prevents', 'inhibits' or 'triggers' – and provides a different counterfactual analysis of each.[9] The problem is that we have one kind of causal claim, one kind of counterfactual. If we know the causal claim, we can assert the corresponding counterfactual; if we know the counterfactual, we can assert the corresponding causal claim. But we never get outside the circle.

The same is true of manipulation accounts. We can read these accounts as theories of what licenses us to assert a causal claim or as theories that license us to infer that when we manipulate a cause, the effect will change. We need a theory that does both at once. Importantly it must do so in a way that is both justified and that we can apply in practice.

[8] Lewis (1973).   [9] Hall and Paul (2003).

6   Hunting Causes – and Using Them

This brings me to the point of writing this book. In studying causality, there are two big jobs that face us now:

> Warrant for use: we need accounts of causality that show how to travel from our evidence to our conclusions. Why is the evidence that we take to be good evidence for our causal claims good evidence for the conclusions we want to draw from these claims? In the case of the two kinds of causal system discussed in part II, it is metaphysics – the theory of probabilistic causality for the first and the assumption of causal priority for the second – that provides a track from method to use. That is the kind of metaphysics we need.

> Let's get concrete: our metaphysics is always too abstract. That is not surprising. I talk here in the introduction loosely about the probabilistic theory of causality and causal priority. But loose talk does not support proofs. For that we need precise notions, like 'the causal Markov condition', 'faithfulness' and 'minimality'. These tell us exactly what a system must be like to license Bayes-nets methods for causal inference and Bayes-nets conclusions. What do these conditions amount to in the real world? Are there even rough identifying features that can give us a clue that a system we want to investigate satisfies these abstract conditions? In the end even the best metaphysics can do no work for us if we do not know how to identify it in the concrete.

By the end of the book I hope the reader will have a good sense of what these jobs amount to and of why they are important. I hope some will want to try to tackle them.

Part I

Plurality in causality

1   Preamble

The title of this part is taken from Maria Carla Galavotti.[1] Galavotti, like me, argues that causation is a highly varied thing. There are, I maintain, a variety of different kinds of relations that we might be pointing to with the label 'cause' and each different kind of relation needs to be matched with the right methods for finding out about it as well as with the right inference rules for how to use our knowledge of it.[2]

Chapter 2, 'Causation: one word, many things', defends my pluralist view of causality and suggests that the different accounts of causality that philosophers and economists offer point to different features that a system of particular causal relations might have, where the relations themselves are more precisely described with thick causal terms – like 'pushes', 'wrinkles', 'smothers', 'cheers up' or 'attracts' – than with the loose, multi-faceted concept *causes*. It concludes with the proposal that labelling a specific set of relations 'causal' in science can serve to classify them under one or another well-known 'causal' scheme, like the Bayes-nets scheme or the 'structural' equations of econometrics, thus warranting all the conclusions about that set of relations appropriate to that scheme.

Whereas ch. 2 endorses an ontological pluralism, ch. 3, 'Causal claims: warranting them and using them', is epistemological. It describes the plurality of methods that can provide warrant for a causal conclusion. It is taken from a talk given at a US National Research Council conference on evidence in the social sciences and for social policy, in response to the drive for the hegemony of the randomized controlled trial. There is a huge emphasis nowadays on evidenced-based policy. That is all to the good. But this is accompanied by a tendency towards a very narrow view of what counts as evidence.

In many areas it is taken for granted that by far the best – and perhaps the only good – kind of evidence for a policy is to run a pilot study, a kind of mini version of the policy, and conduct a randomized controlled trial to evaluate the

---

[1] Galavotti (2005).
[2] See Cat (forthcoming) for a discussion of different kinds of causality which apply to diverse cases in the natural, social and medical sciences.

effectiveness of the policy in the pilot situation. All other kinds of evidence tend to be ignored, including what might be a great deal of evidence that suggested the policy in the first place.

This is reminiscent of a flaw in reasoning that Daniel Kahneman and Amos Tversky[3] famously accuse us all of commonly making, the neglect of base rate probabilities in calculating the posterior probability of an event. We focus, they claim, on the conditional probability of the event and neglect to weigh in the prior probability of the event based on all our other evidence. It is particularly unfortunate in studies of social policy because of the well-known difficulties that face the randomized controlled trial at all stages, like the problem of operational-izing and measuring the desired outcome, the comparability of the treatment and control groups, pre-selection, the effect of having some policy at all, the effects of the way the policy is implemented, the similarity of the pilot situation to the larger target situation and so on.

Chapter 3 is so intent on stressing the plurality of methods for claims of causality and effectiveness that it neglects the ontological pluralism argued for in ch. 2. This neglect is remedied in ch. 4. If we study a variety of different kinds of causal relations in our sciences then we face the task of ensuring that the methods we use on a given occasion are appropriate to the kind of relation we are trying to establish and that the inferences we intend to draw once the causal claims are established are warranted for that kind of relation. This is just what we could hope our theories of causality would do for us. 'Where is the theory in our "theories" of causality?' suggests that they fail at this. This leaves us with a huge question about the joint project of hunting and using causes: what is it about our methods for causal inference that warrants the uses to which we intend to put our causal results?

[3] Kahneman and Tversky (1979).

2   Causation: one word, many things

2.1   Introduction

I am going to describe here a three-year project on causality under way at the London School of Economics (LSE) funded by the British Arts and Humanities Research Board. The central idea behind my contribution to the project is Elizabeth Anscombe's.[1] My work thus shares a lot in common with that of Peter Machamer, Lindley Darden and Carl Craver, which is also discussed at these Philosophy of Science Association meetings. My basic point of view is adumbrated in my 1999 book *The Dappled World*:[2]

The book takes its title from a poem by Gerard Manley Hopkins. Hopkins was a follower of Duns Scotus. So too am I. I stress the particular over the universal and what is plotted and pieced over what lies in one gigantic plane . . .

About causation I argue . . . there is a great variety of different kinds of causes and that even causes of the same kind can operate in different ways . . .

The term 'cause' is highly unspecific. It commits us to nothing about the kind of causality involved nor about how the causes operate. Recognizing this should make us more cautious about investing in the quest for universal methods for causal inference.

The defence of these claims proceeds in three stages.
    Stage 1: as a start I shall outline troubles we face in taking any of the dominant accounts now on offer as providing universal accounts of causal laws:[3]
1 the probabilistic theory of causality (Patrick Suppes) and consequent Bayes-nets methods of causal inference (Wolfgang Spohn, Judea Pearl, Clark Glymour);
2 modularity accounts (Pearl, James Woodward, economist Stephen LeRoy);
3 the invariance account (Woodward, economist/philosopher Kevin Hoover);
4 natural experiments (Herbert Simon, Nancy Cartwright);

---

[1] Anscombe (1993 [1971]).   [2] Cartwright (1999), ch. 5.
[3] I exclude the counterfactual analysis of causation from consideration here because it is most plausibly offered as an account of singular causation. At any rate, the difficulties that the account faces are well known.