Erkenn
DOI 10.1007/s10670-016-9869-8

ORIGINAL RESEARCH

Big Systems Versus Stocky Tangles: It Can Matter to the Details

Nancy Cartwright[1]

Received: 18 December 2016 / Accepted: 18 December 2016

Abstract Wolfgang Spohn's Frege prize lecture, like the work on which it is based, is a tour de force of rich, elegant, coherent argument about how the projected world that we experience is constructed. But we do not live in this projected world nor reason about it. The things Spohn constructs are there from the start—or so my Stanford School pragmatism teaches. This paper explores a deep difference in philosophical approaches—Spohn's elegant proofs versus the stocky, tangled arguments I advocate—and illustrates how these play out in far more detailed disputes about the nature of causality and causal inference.

1 Introduction

Wolfgang Spohn and I have discussed causation for a long time, very often disagreeing on the details: do causes increase the probability of their effects, does the causal Markov principle hold for all 'complete' sets of causal relations, is there one concept of 'cause' or many, are there only as many causes as needed to reflect the probabilities or many more, does the concept of direct causation make sense, and more. We also have very different philosophical viewpoints and concomitant methodologies. What I had not realised till working on this commentary on Spohn's so-well-deserved Frege Prize lecture is how much the little, local disagreements are connected with the big differences.

Spohn's work is a tour de force, an impressive edifice of formal argument to produce a model world. I by contrast follow a piecemeal approach. What I hope to do here is to show how these big-picture differences can underpin real differences in the philosophical details, using causality to illustrate. This is not in aid of arguing

✉ Nancy Cartwright
nancy.cartwright@durham.ac.uk

1 Durham, UK

Published online: 19 January 2017

for one approach or the other, nor for one stand or another on this or that issue of causality but rather to see better how local disagreements converge with broad methodological differences.

I am an advocate of the Stanford School, which I characterize by: pluralism, particularism and concern with practice. The Stanford School is epitomized for me by Patrick Suppes' demand: "Let's get down to the details." I couple that with Suppes'—and Spohn's—insistence on rigor. But Spohn and I differ on where the rigor lies. My particularism bans idealisations that make things look the same only by not looking very carefully, and the commitment to practice keeps my arguments low to the ground. Also, consonant with Stanford School particularism and pluralism, comes scepticism of "the big system", which marks a clear divide between Spohn and me. What I worry about *the big system* is reliance on long arguments and lofty generalisations.

In a *long argument* with many separate premises, the conclusion fails if any of its premises fail. No step in one of these argument is likely to be entirely uncontroversial; each is at least a little dicey. The probability, then, that one or another, somewhere along the line, is false is significant. So why should we trust the conclusion, even if we cannot spot where the error lies? It is on account of this concern that I oppose

Tall, skinny arguments that are sparse and tidy

In favour of arguments that are

Short, stocky, and tangled.

'Stocky' is wide—the arguments cover a lot of the territory under the conclusion—and solid. I use 'tangled' to describe a rich network of interrelated arguments, each firmly attached to the ground, some with shared premises but where a great many also have a number of independent premises and, importantly, premises that come from a variety of different places outside the immediate domain in which the conclusion lives. Long, tall arguments can be beautiful, elegant, an intellectual triumph. But where security of conclusions matters, we had better stick to arguments that are short, stocky, and tangled.

1.1 Lofty Generalisations

My view about lofty principle is much like Pierre Duhem's of the laws of physics. They are rough templates; they organize a great deal of material under them—so long as you don't look too closely at the details and you avert your gaze from items that you might have wanted to cover in the first place but that don't fit into the box. Lofty principles do an essential job of organization for us; and they are often great tools for beginning to build accurate models. But when it is security we want in our conclusions we should look far closer to the ground.

By contrast with the construction of the big system, this piecemeal particularism may look very unphilosophoical. For instance, Pierre Duhem tells us about some of the British ancestors of this tradition:

> If the mind of Descartes seems to haunt French philosophy, the imaginative faculty of Bacon, with its taste for the concrete and practical, its ignorance and dislike of abstraction and deduction, seems to have passed into the life-blood of English philosophy. "One by one, Locke, Hume, Bentham, and the two Mills[…]" All these thinkers proceed not so much by a consecutive line of reasoning as by piling up of examples. Instead of linking up syllogism, they accumulate facts. (Duhem 1991, 67)

Locke, Hume, Bentham and the two Mills: these are paradigms of Suppes' injunction to get down to the details.

Let us look now at issues of causality to see how these general philosophical differences can crop up again and again, shaping concrete philosophical conclusions at a local level.

2 Getting Hold of Causes

Let us call the world of objects in states of affairs that Spohn constructs—i.e. his 'totality of facts'—the 'projected world'. Independent of whether there are unprojected objects doing things and independent of whether Spohn's epistemically-based pragmatic theory of truth can save him from idealism, I do not agree with Spohn about the status of causation in his projected world.

Spohn undertakes a grand Kantian construction to end up with a Humean world —a world where causation is literally 'not to be seen'. It is not there among Spohn's states of affairs, let alone among those states of affairs that are facts. (Facts recall are the states of affairs we arrive at in the ideal limit of inquiry.) That's why he needs the excursion into reasons to get causality and other, as he terms them, 'modal facts in the ordinary sense, i.e. …the so-called natural modalities' (Spohn 2016, 13).

Everything but the one concrete universe is for Spohn an epistemic projection, but causation is an epistemic projection within an epistemic projection. There are objects and states of affairs within Spohn's projected world, but which kinds of states of affairs? His is a strange under-populated world, devoid of all causings as such. There are no pushings; no pullings; no teachings or learnings; no smotherings or uncoverings; no eliminatings or restorings; no gratings, choppings, bakings, whiskings, sauteeings, boilings; no beheadings, invadings, executings, enslavings or freeings; no helpings nor hinderings; elecctings nor just taking charges,… These are all causings in the general sense proposed by GEM Anscombe, that a cause produces or is responsible for its effect (Anscombe 1971); the cause makes the effect happen; the effect comes out of the cause, which is, I would argue, what is in common among the various happenings that we want to label 'causal' for various purposes in various settings. They are all causings and no causings exist in this projected world.

We may, as Spohn allows, assign a truth value—indeed the truth value 'true'—to 'Nancy caused the boulder to fall'. But not because any actual pushing by me can be found in our projected world. Rather because Nancy [but what feature of Nancy if not her pushing?] is a reason for the falling of the boulder, a conditional reason 'given the entire history [of that world] up to' (Spohn 2016, 16) the falling of the boulder. Nor does causality supervene on some complex of facts. In particular, although Spohn is responsible for significant advances in work relating causality to networks of facts about probabilistic dependencies and independencies, causal claims are not summaries of any complex of facts about objective probability relations.

Spohn is surprisingly like Bas van Fraassen, who also puts causal facts in a special category, a step beyond even theoretical entities. Spohn trades eventually in ontology whereas van Fraassen's concerns are consistently epistemological, with what rationality permits and compels belief in, and correlatively with what knowledge science should provide us with. But both arrive at a 3-tier society, each tier less respectable than the one before, with my poor friend, Ansombian causality, relegated to the outermost circle.

This is a new level of epistemic projection. As Spohn says, "It is entirely subjective; what speaks for what is determined by my subjective epistemic state or, more specifically, by my conditional degrees of belief" (Spohn 2016, 15). Spohn tells us, "Since reasons are relative to an epistemic state, causation is so too…" But, as he urges, "causal judgements involve very special beliefs: given the entire history up to *t,* what do I believe to happen at *t*" (Spohn 2016, 17).

This relativisation to an entire history is important. It is an essential ingredient in Spohn's escape from idealism. The other essential ingredient involves laying down some principles about the nature of causal reasoning, Spohn's 'objectivization conditions of causal beliefs'. Once these assumptions are in place, he can show that his dynamic theory of inductive inference—his theory about how reasons for changing one's beliefs work—becomes a theory of causal inference. He tells us in *Laws of Belief:*

> The causal relation is just the objectifiable part of our much richer and more disorderly reason relation. In other words, if we want to objectify our inductive strategies, if we want to align our dynamics of belief to the real world, we have to attend to causation, to the objectifiable part of our reasons. This is what the notion of causation is for. (Spohn 2012, 469)

Let us look more closely at this reasons relativisation. It is familiar in epistemology to suppose that reasons are reasons for agents each with their own peculiar background beliefs. But Spohn needs something stronger than the usual relativity of reasons to background beliefs. For his purposes the verdict about reasons must depend on a specific set of background beliefs: those about the entire history of the world. This calls into play a very special agent, one with beliefs about that entire history. But it is not enough that the agent have beliefs about what the entire history of the world is, those must be *true.* If not, Spohn's truth conditions give the wrong results. So I think we need to read Spohn thus:

> SC: x causes y iff x is a reason for y for an agent who believes H is the history of the world up to the time of y AND H is the history of the world up to that time.[1]

This may seem like minutiae, but it matters when it comes to real, on-the-ground, practical epistemology: how can we actually go about finding out about the world? In judging that x causes y we must, à la Spohn, judge whether x would be a reason for y given that one knows the entire history of the world. That's a tall order! To be fair, it is not as tall as may at first appear. For this does not require one to know the entire history but merely to judge what would be reasons if one did. That may not be at all impossible.

By now most philosophers will be familiar from the work of Gerd Gigerenzer and others with the idea of a 'cheap heuristic'[2]—a method of thinking through a problem that gets one to the right end point (or near enough often enough) without going through a laborious 'brute force' calculation. Moreover we know we don't need to look at the entire history of the world to make a reliable causal judgement. Our standard comparative methods in the sciences look for differences between the presence and absence of causes controlling for a far smaller domain of features[3]: the other causes operating to produce y at the time. If x is a reason for y once we know those, then x is a cause of y. So we call into play Principle CC, which I present in its crudest form for simplicity of discussion[4]:

> CC: x causes y iff x is a reason for y for an agent who knows all the other causes influencing y.

CC, I claim, is the lesson that probabilistic theory of causality teaches and it is the justification that lies behind our standard scientific methods for causal inference.[5]

Why does Spohn not use CC rather than SC from the start? The answer for conventional Humeans is that they want to offer a reductive account of causation: necessary and sufficient conditions for 'x causes y' that refer only to facts in a 'Hume world', i.e. a world where states of affairs consist of objects exhibiting some special kind of non-modal properties. CC judges whether x causes y relative to other

1 An anonymous referee underlines that it is not the agent who knows the whole history who is to judge that x is a reason for y but 'us theoreticians'.

2 Cf Gerd Gigerenzer et al., *Simple Heuristics that Make Us Smart* (New York: Oxford University Press, 1999).

3 Consider for example a later replication of a physics experiment to test a causal hypothesis. What matters is that the other causes for the effect are the same, or accounted for, not any further information about the history of the world. This is also assumed in our use of statistical methods for testing causal claims, like stratification (we stratify on other causes and take that to be sufficient if it could be done), randomised controlled trials, qualitative comparative analysis, causal-Bayes-nets methods, and instrumental variables and other econometric methods (in these latter the 'exclusion' conditions refer to other causes, not to arbitrary factors in past history). The assumption is also clear in JL Mackie's (1974) famous and widely adopted INUS theory of causality (at least the claim that causes are INUS conditions is widely adopted though not the converse).

4 Both Spohn and I adopt something along these lines. Not surprisingly, our two versions are different.

5 Cf Cartwright (2007), chapter 3, 'Causal claims: Warranting them and using them'.

causes of y. The Humean can avoid relativising to other causes by relativising to the whole history of the world, which is bound to catch all other causes.[6] Spohn's motivation is similar:

> The crushing conclusion [from CC] is clear: any explanation of the circumstances must refer to the notion of causation. The circle is vicious… We stand at a crossroads here. Should we continues to strive for an explication of the notion of causation that results in an explicit [reductive] definition?…I am heading for the more ambitious project. (Spohn 2012, 355)

Things may not be so bad between CC and the Humeans as they first appear however. CC constrains labelling the relation between x and y 'causal' by how the relations to y of other features are labelled—are these labelled 'causal' or not? But CC equally constrains whether the relation between any other feature, say z, and y is labelled 'causal' by whether the relation between x and y is labelled 'causal'. CC thus acts as a consistency constraint on the full set of relations labelled 'causal' with respect to y.

Consistency constraints can be a powerful tool, as we see from their use in various sciences. For instance, the first serious quantum theory of the laser (by Willis Lamb) relied on a 'self-consistency' model. Perhaps more familiar are rational expectations models in economics, which demand that the mean values that agents in the model predict for certain quantities (i.e. the agents' expectations in the sense of what the agents anticipate will happen) are the same as the mean values for those quantities that result from the agents' actions (i.e. the expectations of those quantities in the statistical sense). This consistency constraint, coupled with other principles in the model, fixes the behaviours of the agents uniquely.

Unfortunately CC by itself is not enough to dictate a unique verdict about what sets of features are causally related. But there are further constraints on the causal relation we could adopt that would do the job. This is what I see Spohn's objectification conditions doing. If causal relations are to satisfy this extended set of constraints then, given our pre-settled reasons relations and the way the world actually is, there is only one verdict possible about whether x causes y for each pair of features $\langle x, y \rangle$.

This however creates a puzzle for me about why Spohn opts for SC rather than CC. He after all (unlike me) is inclined to accept his objectification conditions—and these, if I understand them correctly, operate essentially on CC not SC. Moreover, adopting SC creates a special demand given that the narrower CC is sufficient. Because now a defense is required that SC implies CC. This involves showing at least three things. First, that the huge amount of extraneous information contained in the entire history of the world never produces a different verdict than the narrower information about what other causes influence y. Second is a consistency condition, that CC for x causes y follows from SC for x causes y supposing that SC is used in CC as the criterion for whether or not each other factor is a cause of y. Third, when SC is used as a criterion for judging which are the 'other causes' to be used in CC, it gets them right.

[6] This supposes that relativising to more than we need does not change the verdict.

Connecting with CC is no idle matter, especially if one starts from a concern with the real epistemic practices that get us around in the world.[7] CC is at the core of our many methods for causal inference that depend on Mill's method of difference, including both 'single-shot' controlled experiments and methods that look for probabilistic differences, as in econometric causal models, stratified observational studies and randomised controlled trials. Consider a 'single-shot' controlled experiment to see if x causes y: we fix, to the best of our ability, the other causes of y, then look to see if the presence versus absence of x makes a difference. Here we have not only epistemological and manipulability problems with employing SC but it is literally impossible for the entire history up to y to be the same in the settings with x and without x.

Supposing we do require that SC imply CC. There are three strategies one can pursue to defend SC. The first is argument by philosophical example. This however is better suited for showing that things are not equivalent than for showing that they are. It also suffers from the well-known problem that intuitions about what verdicts are correct in various cases differ and can depend heavily on the context to which a case is assigned in the process of evaluation.

The second strategy is to establish an empirical connection. Here we face the general problem that Hasok Chang labels 'the experimenter's regress' (Chang 2004). We want to see whether A, which we know how to find out about, is a good indicator of B, which we don't: is CC a good indicator of SC? But we can't establish a correlation because we have no way to figure out when B holds in the first place. This is a standard problem in developing measurement procedures in the sciences where we use a variety of methods together in an ongoing feedback process to solve it. For instance, where there's some reasonable starting theory of the sources or the consequences of B, we look to see if those occur in the face of A. We also look for convergence among the verdicts of different indicators each of which has some starting plausibility. But these strategies are not so promising for the case of SC and CC.

First, because in the sciences we often evolve our notion of A—for example, just what temperature is—in tandem with refining the details of the Bs that are supposed to be indicators of it. This wouldn't serve to defend SC as it is, and modifying SC would threaten the impressive network of proofs that underwrites Spohn's account of causality.

Second, I don't see how to carry it off. What corresponds to the empirically testable sources or consequences of the set of factors judged relevant by SC that might correlate with CC's relevance set? Or, what other methods are plausible starting methods for identifying the SC relevance set that might converge in their verdicts with the CC method for doing so? We can agree, I believe, that there are many cases where a variety of different kinds of indicators that x caused y converge, and converge with the verdict of CC (indicators like the character of the effect [does y occur at the time, in the manner and of the size to be expected had x caused it?], various symptoms of causation [e.g. side effects that could be expected had

[7] This discussion is foreshadowed in footnote 1, which was added in response to a call by an anonymous referee to expand at that point.

x operated to produce y], the presence of requisite support factors [moderator variables] and absence of inhibitors, and the presence of expectable intermediate steps [mediator variables]). The problem is to provide some reasonably good starting argument that these are connected with SC.

The third method is proof, and here Spohn has something incredible to offer: a series of proofs that "instead of the largest possible circumstances… [NC: i.e. the entire history of the world] we may as well base our explication of direct causation… on much smaller [sets of] circumstances … and thereby arrive at exactly the same direct causal relations" (Spohn 2012, 373). Here different smaller sets will recoup slightly different kinds of direct causes. None of these are exactly CC, which Spohn does not believe is always correct,[8] but under one special assumption that he thinks can often hold, SC causal verdicts will be the same as those of CC.

Why am I still worried? Because I don't see how to satisfy my third desideratum: that we have reason to think the causes picked by SC are correct. How can I doubt that, given that, for purposes of this discussion, I have assumed that CC gets it right? My reason lies in one of the strengths of the proof itself: it does just what I said is demanded for consistency, which is to show that SC implies CC, where the 'other causes' referred to in CC satisfy SC. But I want them to satisfy CC and it seems we need this very proof to show they do. This is perhaps fixable by some kind of expanded proof by induction. So I move to my more general worry, which returns me to the overall theme of short and stocky versus tall and thin.

The argument is too long—way too long, and many of the premises are controversial. Let's start at the end of the chain. We don't get a unique verdict from SC about whether x causes y without additional assumptions. The central ones Spohn employs are close relatives of the three central axioms characterising causal Bayes nets. Quite independently of Spohn's use of them for objectification, I have quarrelled with every one of them… and for separate reasons.[9] For each there are, I maintain, many sets of relations that we wish to label 'causal' that do not satisfy that axiom. Not only do we wish to apply that label to these relations but doing so does real work for us of one kind or another that we expect of causal relations, like helping to assign responsibility and blame, or to predict what will happen, to understand the world better, or to change it.

I do not suggest that I and others are right in all our criticisms of all these Bayes nets axioms. The problem is that they make for a very tall argument for objectification. All these objections and counterexamples could be wrong. But that seems extremely unlikely; and, if they are not, there's a reasonable chance that one

[8] Though for different reasons than mine for modifying CC. Mine depend first on the fact that there are a variety of other reasons than common causes why two factors may be correlated, whereas CC deals at base only with common causes (e.g. x may be a reason for y even once common causes have been taken into account if they are produced in tandem by a probabilistic cause; or if they are joint causes of the same effect in a population where that effect is universally present; or if they change in the same direction in time); second, because x may cause y but not be a reason for it, for instance where its positive effects on y and its negative balance. Cf Cartwright (2007), chapter 6, 'What is wrong with Bayes nets?'.

[9] Cf Nancy Cartwright, *Hunting Causes and Using Them: Approaches in Philosophy and Economics* (Cambridge: Cambridge University Press 2007) and *The Dappled World: A Study of the Boundaries of Science* (Cambridge: Cambridge University Press 1999).

of these premises is not a universal condition on causation. If so, obectifivization is out the window unless or until other arguments are to hand, and I would want lots of other arguments to make sure that the whole structure far more stocky.

Looking back further, there is Spohn's theory of reasons to buy since assumptions about how reasons operate are essential for the theory of causation, both for the proof that SC implies versions of CC and for the proof of uniqueness that turns epistemic causes into objective ones. Here again one can have quarrels with specific assumptions but what I would like to draw attention to is a kind of relativisation in the theory of reasons that endangers the objectification of causation: the relativity to choice of model. Although the relativity of causes to reasons may be eliminated by a successful uniqueness proof, no such proof is in the offing for model-relativity, as Spohn himself notes: cases of causality are always described by a "(conceptual) frame" that specifies the set of variables to be used (Spohn 2012, 341). Ranks and probabilities, which formalise reasons for Spohn, both require event spaces to range over, as does the notion of the entire history of the world used in defining direct causation. Further, the concept of direct causation requires dividing time into discrete chunks so that nothing can get in between a direct cause and its effect. Verdicts relative to one formal frame can differ from those relative to another, as Spohn stresses in his 1990 paper, 'Direct and Indirect Causes'. This paper relies on probabilities rather than the ranks of his 2012 *Laws of Belief* but the treatments of causality in the two are parallel. Spohn explains:

> … everything said about causation is relative to the descriptive frame given by the [event space]…[This relativization] is essential because the causal relations may indeed vary with the frame. (Spohn 1990, 125)

When this can happen, the objectification constraints can pick only a set of causal relations that are unique up to the choice of frame, which is no objectification at all.

Some frame relativity can be shown to be harmless. Sometimes it doesn't matter whether we fine-grain or coarse-grain the event space and sometimes expansions and contractions do not matter. But these kinds of results are not sufficient to eliminate frame relativity.

Perhaps, though, following ideas from Spohn's pragmatic theory of truth on the limits of scientific exploration, "where absolutely everything is explored" (Spohn 2016, 12), we can hope that the self-correcting methods of science will tend towards a unique choice. After all, Spohn tells us, "In the limit of inquiry not only our beliefs are guaranteed to be true; we also know what the states of affairs and what the facts are" (Spohn 2016, 13). Spohn's work on the dynamic laws of belief has certainly made impressive inroads on the formal and philosophical side to, as he puts it, "develop [the pragmatic theory of truth] so as to meet theoretic standards" (Spohn 2016, 12).

The problem is that I find little in the practices of the sciences to support this hope. For me the job is not to make sense of the (in Spohn's terms) "nice metaphoric" (Spohn 2016, 12) of 'the limits of inquiry' but rather to see that such limits make sense for real inquiry into the real world. When we get down to the details, as Suppes urged—the details of practice—rather than convergence I see ever more unrelated proliferations of event spaces from different subdisciplines looking at not only different problems and different aspects of them but often looking at what appears to be the same aspect of the same problem but with a different methodological slant.

Of course the limit is traditionally described as the 'ideal' limit to abstract from the details of the real. That's okay so long as *idealisation* is what Catherine Elgin calls a 'felicitous falsehood' (Elgin 2012). A felicitous falsehood puffs up the truth about features of interest and lies about many others in order to highlight the characteristics we care about and eliminate ones that mask them from us. But it has its feet firmly in the actual; it cleans up what is there, presents it in a better light. It does not whisk away the actual and drop something more desirable in its stead. This latter kind of idealization, which is closer to what we see in Spohn, does not support Spohn's project of objectivization.

But perhaps Spohn would insist it does. For in discussing speculation about "general, or the most general, laws of belief" that are "at best loosely connected with empirical facts", he claims: "… it is certainly desirable to conceive the manifold of relevant phenomena under some general laws, even if only as idealizations that apply to reality only with massive help from correcting theories" (Spohn 2012, 7).

This echoes my own view, and that of many instrumentalists and pragmatisists (and notably Duhem whom I have quoted) about the general laws of high physics theories. Conceiving the phenomena under these laws, even if they only apply with massive corrections (which, I have argued, practice shows is the only way they apply), is indeed desirable, and for a host of reasons. The general laws provide an all-purpose tool that, when used with the right other tools in the right ways, case-by-case, can help in the construction of close-to-the-ground models that apply with hardly any corrections. And, as Mary Morgan points out in her work on the travel of facts and techniques from one domain to another, seeing that the more low-level laws grouped together under the same general claim are similar in significant ways allows us to use similar methods of study, modelling strategies, approximation techniques and the like and it suggests analogous predictions to look for from one domain to another (Morgan 2009).

So in the sciences, idealisations that fit the world only with massive corrections can be useful for many purposes. But I don't see how this kind of idealisation can serve Spohn's aims, either with respect to objectifying laws of nature or for the project under discussion here of achieving an appropriate event space for the objectivization of causes. If we are to base things that matter about the world on what is arrived at in the limits of scientific exploration, "where absolutely everything is explored", should that not be in the limits of exploration as we actually engage in it when we are successful in finding out about the world—or at least seem to be successful judged by our most rigorous criteria?

To support convergence in the face of my practice-based worries, one might turn to our successful 'high' theories in physics or perhaps biology. I say 'high' theory because that is where we might expect both to overcome proliferation and to find convergence. Can the concepts these theories employ provide the requisite event space and can this be expected to converge on some limit in the end of inquiry? The answer I think is 'No' to both. Consider convergence first. Physics high theory has an exciting history, a monument to the human imagination in part because it has exhibited such dramatic changes in the kinds of features it introduces. That is an old theme from the history of science, much discussed, especially in the US, in the 'science wars' of the 1990s. So I'll say no more about it here.

What about the event space itself? Can it overcome the problems of proliferation? No: the very virtue that leads us to turn to it makes it unfit for the jobs Spohn undertakes. For the event space supplied by high physics theory will not include the bulk of the actual causes that make things happen in the world. I note three strands of practice-based argument in support of this pessimistic conclusion.

The first is short. Physics theory consists of equations and there are no causes to be found in these equations. I myself have argued against this, as have Mathias Frisch and others (Frisch 2012). But it seems we are in the minority. I won't rehearse the arguments on both sides. I just note that in case Frisch and I are wrong, this presents a real challenge to the hope to turn to physics theory for an event space in which causes can be objectified.

The second argument is one I have developed in a paper in memory of Lorenz Krueger (Cartwright 1998). When it comes to producing outcomes in the real world, or describing how they come about, even ones where our knowledge of high physics theory plays a central role, we utilise mixed, untidy collections of concepts from a wide range of sources, including engineering, ordinary good sense and hands-on tacit knowledge of how things work together. Even the physics comes from many different, narrow, unorchestrated special subfields. Physics may be the Queen of the disciplines but she does not dictate what happens; she is only part of a motley assembly. We can try to idealise away the motley assembly but then we lose the ability to predict and control. Or we can assume that the ways we produce and explain outcomes will be very different in the future than they have been in the past. But I see no grounds for betting on this.

The third argument has to do with the way physics concepts hook onto the world. Here is one picture. The concepts used in producing and explaining real outcomes only look motley because we are looking at too concrete a level. It is well known that things that look different in the concrete may be the same in the abstract. The trajectory of a coin falling to the ground, of a planet circulating the sun and of a cannonball shot over a moat may look different at that level of description but they are in fact all motions subject to a central gravitational force. If this view is a plausible idealisation (in the good sense) of what is going on, then maybe the single event space can be salvaged.

But I don't see any way to defend it in the face of the empirical record. My particularism here follows two heroes of mine, Otto Neurath and JL Austin. Neurath pointed out that the exact concepts of proper science do not fit the blousy empirical world in which we conduct the business of producing outcomes and explaining why things happen. For most of the concepts we use in causal reasoning, we have no idea how to relate them to those of physics. Even for those we connect with physics, it takes a lot of additions and subtractions, a great many of which are ad hoc. They are there just to bend the concepts of physics to fit the phenomena. This is part of the reason that the laws of physics fit onto the world only with massive corrections and it was Neurath's reason for rejecting the project of theory of confirmation to substitute the loose notion of shaking. Empirical results cannot confirm or disconfirm theory in actuality; they can only shake or steady our confidence in it.[10]

Austin was similarly suspicious of the attempt to sweep everything into one tidy space:

> This is by way of a warning in philosophy. It seems too readily assumed that if we can only discover the true meanings of each of a cluster of key terms…that we use in some particular field (as, for example, "right", "good" and the rest in morals), then it must without question transpire that each will fit into place in some single, interlocking, consistent conceptual scheme. Not only is there no reason to assume this, but all historical probability is against it…. We may cheerfully use, and with weight, terms which are not so much head-on incompatible as simply disparate, which just don't fit in or even on. Just as we cheerfully subscribe to, or have the grace to be torn between, simply disparate ideals—why must there be a conceivable amalgam, the Good Life for Man? (Austin 1956–1957, 29, note 16)

I have here no new arguments for the proliferation of concepts necessary to do causal reasoning. That is the meat of my book *The Dappled World.* But I had not before seen how those views, grounded in the demand to 'get down to the details' and to look at the practices of knowledge acquisition and use in contrast to the promissory notes, intersects so consistently with Spohn's project. The difference in basic philosophical method seems reflected in specific disagreements at many unexpected internal points.

3 No Causes In, No Reasons Out

I focus on causes not only because it is where Spohn and I have met most often but because I fear that the failure of causes to appear in Spohn's projected world from the start undermines his plan to procure them via reasons. For, I urge, there are no reasons without causes, or more accurately, without some forms of what Spohn calls 'natural modalities', like causes, laws, powers, or perhaps Stephen Mumford's dispositional modality.

Why? A reason for y must, as Spohn says, 'speak for' y. But it must also be able to speak *to us*. This latter causes trouble. But let's start with the former. This is what I call 'evidence' in my theory of evidence for policy and where-soever rigour matters.[11] Whether (the fact referred to by) x speaks for (that referred to by[12]) y or not depends on what else is true. That is why I picture evidence as a 3-place relation: x speaks for (is evidence for) y relative to argument A iff A is a valid argument for y with true premises and x is a premise in A that is essential for the

[10] Cf Nancy Cartwright et al., *Otto Neurath: Philosophy Between Science and Politics,* (Cambridge: Cambridge University Press 1996).

[11] See Nancy Cartwright, *Evidence: For Policy and Wheresoever Rigor is a Must.* (London: LSE 2013) and Nancy Cartwright and Jeremy Hardie, *Evidence Based Policy: A Practical Guide to Doing it Better.* (New York: Oxford University Press 2012).

[12] For simplicity of presentation I shall hereafter drop attempts to avoid use/mention conflations.

derivation of y. This a very non-restrictive notion of evidence since there are always valid arguments that turn any actually occurring fact into evidence for any other. If x and y are both true then so is $x \supset y$ and thus x is evidence for y relative to an argument for y with x and $x \supset y$ as premises.

I don't think this is harmful for a general account of evidence. x does speak for y relative to the truth of $x \supset y$. If you know $x \supset y$, then surely finding out that x holds is a good way to find out that y does. But can x thereby be a reason for us for y? That depends (at least) on whether $x \supset y$ could be available to us. We can learn x is true but that is no help with respect to y if the only way to access $x \supset y$ is by learning y. There are facts about the world that will do the job, facts that—if we can learn them at all—we can learn without directly learning y. And I would argue that we can learn them since we are never in possession only of data in Spohn's sense, though that is not necessary to my point here. For if we cannot, then even these facts cannot salvage reasons. What matters is that these all involve some one or another of the natural modalities. For instance, if there is a law that y follows x, or it is true that x causes y, or that they share a common cause, or that x and y are jointly sufficient and separately necessary causes of some further fact z that obtains, or x is 'disposed' to y in Steven Mumford and Rani Lill Anjum's dispositional modality sense (Mumford and Anjum 2011), then there is some truth about the world that could supply access for us to $x \supset y$ without having to learn y. But not otherwise. And that seems a real problem for Spohn's programme.

Spohn wants to build what he calls 'an inductive scheme': a "scheme that projects from each possible sequence of data some set of beliefs…"; and he wants to do so in the 'normative mode', in which "we want to know which inductive scheme is the right one" (Spohn 2012, 3). One scheme we know is not the right one is induction by simple enumeration: Swan 1 is white, swan 2 is white, swan 3 is white…So the swans in Sydney Harbour are white. But every scheme that relies only on patterns in a sequence of data is essentially induction by simple enumeration.

We can look to the work of John Norton to support my concerns about inductive schemes. Norton has a two-pronged argument. First he raises problems for the most plausible schemes he can identify. Second he describes what he thinks is actually going on in real life scientific inductions in his 'material theory of induction': "…inductive inferences are not warranted formally but materially. They are justified by facts…" (Norton, manuscript 2014, 2). And the material facts that do this job turn out to be modal.

We can get a good idea what this involves from what Norton calls 'the clearest illustration', Marie Curie's induction on radium chloride:

> The attempt to explicate it with the schema of enumerative induction failed. We could not justify why the schema should be limited precisely to the few properties of radium chloride that Curie so confidently generalized. The justification for this restriction cannot be found in any formal analysis of predicates and properties. Rather, it lies in the researches of chemists in the nineteenth century. The core result is known as "Haüy's Principle" … [which] asserts that, generally, each crystalline substance has a single characteristic

> crystallographic form. The principle is grounded in extensive researches into the chemical composition of crystalline structures and into how their atoms may be packed into lattices. That means that once one has found the characteristic crystallographic form of some sample of a substance, generally one knows it for all samples.
>
> Curie's inductive inference is warranted by Haüy's Principle and not by conformity to any inductive inference schema. There is an inductive risk taken in this conclusion, as indicated by the "generally" in the principle. (Norton, manuscript 2014, 10)

The first thing to note is that although, as Norton indicates, there is inductive risk—Haüy's principle may have exceptions, what matters is that Haüy's principle is just that—a principle. Exceptions will have explanations. Second to note is that Norton takes Hauy's Principle to state a material fact. Third, that it is a natural modality. The principle does not just describe some pattern that has instances in the data. If it did, the inference would after all be in accord with familiar inductive schemes. It depends rather on the nature of radium chloride, on the nature of other chlorides and on further chemical principles:

> We now know that all crystals fall into one of seven crystallographic families…Discerning these families constituted a major mathematical challenge and it was only after the mathematical problem was solved that truly reliable inductive inferences on crystalline forms were possible. When Curie identified radium chloride crystals as just like those of barium chloride, she was adopting the expediency of not specifying the family formally, but merely allowing that it was the same as that of barium chloride. This in turn lent credence to her induction since another principle of chemistry, the law of isomorphism, allowed that analogous chemicals formed similar crystals. (Norton, manuscript 2014, 10, 11)

As in this case, in all of Norton's examples the material facts that support inductive inference turn out to include natural modalities, like 'Generally, each crystalline substance has a single characteristic crystallographic form' or 'Electrons all have the same size charge'. So in Norton's scheme, as in mine, there are no reasons without modalities.

Is this not just Hume's problem of induction then? If so, this should not assuage the worries I have been raising. Hume's argument is very clever. He like Spohn constructs a Hume world, a world without any of the natural modalities. Then he formulates a problem about reasons that cannot be solved without modalities. That problem is a problem for us only if the Hume world is a good model of ours, or, to be more fair to Hume, of our world as we can experience it. But the Hume world is not a good model for ours and there are no compelling arguments that it is once Hume's doctrines about what impressions are, his copy theory of ideas and his accompanying associationist theory of concept formation are rejected, as I take it they should be.

Our world, even as we experience it, is rife with modalities. One might argue, following Wilfrid Sellars, that if we can know that objects have specific properties, we must already have access to some natural laws. That's because properties—like being a radium or a barium chloride crystal—come in families with natural behaviours and natural relations to other properties. We must recognize these if we are to know that we have radium chloride crystals in front of us, even if like Marie Curie, what we know is but a crude version of the natural connections so we are not able to "specify the family formally". A similar conclusion is suggested by the now popular metaphysical doctrine that properties are collections of powers, so that, in whatever sense one knows that an object has a property, one must know that it has at least some of the salient powers that constitute that property.

These are both deep and controversial issues about topics that I do not study in detail. I touch on them to point out that the modality close to my heart, causality, is not alone in its claim to be a part of Spohn's world of facts. There are independent arguments for other modalities to be there as well. My own favourite is causality, which brings me to the question of my last section: Where did all the causes go in Spohn's projected world?

Before that though I should consider one way to try to solve the problem of how to have reasons without first having causes yet in the same breath maintain that we cannot have causes without reasons. That is: reasons (i.e. good reasons) and causes are co-constituted. They are built together. This is an exciting prospect, and a project of a kind that Spohn is a virtuoso at. We are not there yet but perhaps he will get us there. Curmudgeonly me of course would probably not be happy even then because of the idealisations I suspect such a proof would involve, "idealizations that apply to reality only with massive help from correcting theories."

4 Where Did All the Causes Go?

Why do we need an account of causality at the point that Spohn proposes one for? Once we are moving in the projected world, why can we not use the simple Tarski formula for its truth conditions, supposing that the right-hand-side is evaluated at the projected world: '*c* cause *e*' is true iff *c* causes *e*, just as Spohn can presumably use this for other relations (like '*a* is larger than *b*' is true iff *a* is larger than *b*)?

Spohn tells us at the start

- "The world is a collection of facts."
- "The universe is just one big concrete object."
- "Universes are concrete objects… whereas totalities of facts are very complex states of affairs (or collection thereof)" (Spohn 2016, 3).

Here is an alternative more attractive to the practical, particularist heart.

- There are a lot of objects. (Notice though that I do not say *a collection* of objects.)

- They are different. Some are big, some little, some negatively charged, some positively, some are electrically neutral, etc. It seems they may differ in more far more ways than we can imagine or that can be counted.[13]
- Things happen, objects change, objects interact. The ball falls to earth, little children grow big, brightly coloured sweaters fade in the wash, electrons attract protons, cats lap up milk.
- To get one big concrete object takes work, for instance, like that described by Scheibe (1991). He claimed there was just one big concrete object. His argument was roughly this. Physics tells us that everything interacts. But under physics laws, if two things interact, they are no longer subject to evolutionary laws separately but only evolve together. So ultimately there is only one thing subject to the laws of evolution: the whole interacting lot. Scheibe's argument may get us to the conclusion that there is just one big object but it won't serve Spohn's purposes since what it shows is that (supposing laws like the current ones in physics are correct), there is only one object *for physics laws to evolve*; the laws can't evolve the different objects separately. It does not get us Spohn's universe as 'one big concrete object',—indeed, the only *concrete* object—from which other objects appear only under projection.

I urge this alternative as the background for my assumption that the world—even if projected—comes much as it looks to us and much as we describe it, with causings right there along with what (I take it) Spohn would include under the label 'data'. Taking the world as it comes until provided with convincing evidence that it is some other way is part of my Stanford School pragmatism, where the evidence should look much like what we would demand of evidence in real practice. If, with the projection of objects and properties, we arrive at a world with states of affairs and facts, what principle allows in facts that fall under the label 'data' but not a huge swathe of the facts that I see around me?

Presumably there are relations among Spohn's properties. Why then the relations '*a* is to the left of *b*', '*a* is bigger than *b*', '*a* is redder than *b*' but not '*a* pushes *b*', '*a* smothers *b*', '*a* restores *b*', '*a* bakes *b*', '*a* helps *b*', etc. These relations are, as Anscombe argued, observable and a great many are measurable, as I argue in *Nature's Capacities and their Measurement* (Cartwright 1994). We can also provide informative characterisations of what they are and describe how they relate to other relations. I do not see how to mount an argument that once we have projected facts we have not already projected causings. So there seems no need to use reasons as a tool to get causes into the (projected) world. This is all to the good since, if my worries are borne out, there would be no reasons available for the job in a world without causings.

[13] Note that I follow Bas van Fraassen's stricture and do not say 'there are properties and objects differ by which they have.' Nor do I say 'properties are collections of objects'.

5 Conclusion

I have focussed on the effects of broad differences in philosophical viewpoint to show how big differences can affect local quarrels. The point is to show what these broad differences can amount to when we get down to the actual philosophical work. Neither Spohn nor I would want to urge that one of these viewpoints should be pursued at the cost of the other. Surely at least this pluralist conclusion is correct: philosophy needs both.

References

Anscombe, G. E. M. (1971). *Causality and determination: An inaugural lecture*. London: Cambridge University Press.

Austin, J. L. (1956–1957). A plea for excuses. *Proceedings of the Aristotelian Society*, *57*, 1–30.

Cartwright, N. (1994). *Nature's capacities and their measurement*. Oxford: Oxford University Press.

Cartwright, N. (1998). How theories relate: Takeovers or partnerships? *Philosophia Naturalis, 35,* 23–34.

Cartwright, N. (1999). *The dappled world: A study of the boundaries of science*. Cambridge: Cambridge University Press.

Cartwright, N. (2007). *Hunting causes and using them: Approaches in philosophy and economics*. Cambridge: Cambridge University Press.

Cartwright, N. (2013). *Evidence: For policy and wheresoever rigor is a must*. London: LSE.

Cartwright, N., Cat, J., Fleck, L., & Uebel, T. E. (1996). *Otto Neurath: Philosophy between science and politics*. Cambridge: Cambridge University Press.

Cartwright, N., & Hardie, J. (2012). *Evidence based policy: A practical guide to doing it better*. New York: Oxford University Press.

Chang, H. (2004). *Inventing temperature: Measurement and scientific progress*. New York: Oxford University Press.

Duhem, P. (1991). *The aim and structure of physical theory*. 1906. Reprint, Princeton: Princeton University Press.

Elgin, C. (2012). Understanding's tethers. In C. Jäger & W. Löffler (Eds.), *Epistemology: Contexts, values, disagreement: Proceedings of the 34th international Ludwig Wittgenstein symposium, Kirchberg am Wechsel, Austria 2011* (pp. 131–146). Berlin: De Gruyter.

Frisch, M. (2012). No place for causes? Causal skepticism in physics. *European Journal for Philosophy of Science, 2*(3), 313–336.

Gigerenzer, G., Todd, P. M., & the ABC Research Group. (1999). *Simple heuristics that make us smart*. New York: Oxford University Press.

Mackie, J. L. (1974). *The cement of the universe*. Oxford: Clarendon Press.

Morgan, M. S. (2009). Traveling facts. In M. S. Morgan & P. Howlett (Eds.), *How well do facts travel? The dissemination of reliable knowledge* (pp. 3–42). Cambridge: Cambridge University Press.

Mumford, S., & Anjum, R. L. (2011). *Getting causes from powers*. Oxford: Oxford University Press.

Norton, J. (2014). A material defense of inductive inference. Manuscript.

Scheibe, E. (1991). Substances, physical systems, and quantum mechanics. In G. Schurz & G. J. W. Dorn (Eds.), *Advances in scientific philosophy, essays in honor of Paul Weingartner* (pp. 215–229). Amsterdam: Rodopi.

Spohn, W. (1990). Direct and indirect causes. *Topoi, 9,* 125–145.

Spohn, W. (2012). *Laws of belief: Ranking theory and its philosophical applications*. Oxford: Oxford University Press.

Spohn, W. (2016). How the Modalities Come into the World. *Erkenntnis*. doi:10.1007/s10670-016-9874-y.