OXFORD JOURNALS
OXFORD UNIVERSITY PRESS

Of Humean Bondage

Author(s): Christopher Hitchcock

Source: *The British Journal for the Philosophy of Science*, Mar., 2003, Vol. 54, No. 1 (Mar., 2003), pp. 1-25

Published by: Oxford University Press on behalf of The British Society for the Philosophy of Science

Stable URL: https://www.jstor.org/stable/3541877

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at https://about.jstor.org/terms

and *Oxford University Press* are collaborating with JSTOR to digitize, preserve and extend access to *The British Journal for the Philosophy of Science*

JSTOR

This content downloaded from
192.58.125.16 on Sat, 07 Jan 2023 03:46:45 UTC
All use subject to https://about.jstor.org/terms

Brit. J. Phil. Sci. **54** (2003), 1–25

Of Humean Bondage

Christopher Hitchcock

ABSTRACT

There are many ways of attaching two objects together: for example, they can be connected, linked, tied or bound together; and the connection, link, tie or bind can be made of chain, rope, or cement. Every one of these binding methods has been used as a metaphor for causation. What is the real significance of these metaphors? They express a commitment to a certain way of thinking about causation, summarized in the following thesis: 'In any concrete situation, there is an objective fact of the matter as to whether two events are in fact bound by the causal relation. It is the aim of philosophical inquiry to analyze this objective relation.' Through a variety of examples, I hope to cast doubt on this seemingly innocuous thesis. The problem lies not with the word 'objective', but with the word 'the'. The goal of a philosophical account of causation should not be to capture *the* causal relation, but rather to capture the many ways in which the events of the world can be bound together.

1 *The metaphors*
2 *Unpacking the metaphors*
3 *Theories of causation*
4 *The two assassins*
5 *The birth control pills*
6 *The smoker-protector gene*
7 *The bicycle thief*
8 *Further examples*
  8.1 *Indeterminism*
  8.2 *Probability-lowering causes*
  8.3 *Parts vs wholes*
  8.4 *Symmetric overdetermination*
  8.5 *Delayers*
  8.6 *Causation by omission*
  8.7 *Double prevention/disconnection*
  8.8 *Preemptive prevention*
  8.9 *Quantitative variables*
9 *Conclusion*

© British Society for the Philosophy of Science 2003

This content downloaded from
192.58.125.16 on Sat, 07 Jan 2023 03:46:45 UTC
All use subject to https://about.jstor.org/terms

1 The metaphors

There are many ways of attaching two objects together: for example, they can be connected, linked, tied or bound together; and the connection, link, tie or bind can be made of chain, rope, or cement. Every one of these means of attachment has been used as a metaphor for causation. Perhaps no philosopher has made greater use of these attachment metaphors than David Hume. He uses them when expressing the importance of cause and effect reasoning:

> 'Tis evident, that all reasonings concerning *matter of fact* are founded on the relation of cause and effect, and that we can never infer the existence of one object from another, unless they be connected together [...]. ([1740/1978, p. 649])

> *Resemblance* [...] *Contiguity* [...] *Causation* [...] these are the only links that bind the parts of the universe together [...] these are the only ties of our thoughts, they are really *to us* the cement of the universe [...]. ([1740/1978, p. 662])

He uses them when articulating the ingredients of our idea of causation:

> [W]hatever objects are consider'd as causes or effects, are *contiguous* [...]. Tho' distant objects may sometimes seem productive of each other, they are commonly found upon examination to be link'd by a chain of causes [...]. ([1739–40], Book I, Part III, Section II)

And he uses them when arguing that the causal relation is inaccessible to *a priori* reason and direct observation:

> Suppose two objects to be presented to us, of which the one is the cause and the other the effect; 'tis plain, that from the simple consideration of one, or both these objects we never shall perceive the tie, by which they are united, or be able to pronounce, that there is a connexion betwixt them. ([1739–40], Book I, Part III, Section XIV)

> The same difficulty occurs in contemplating the operations of mind on body; where we observe the motion of the latter to follow upon the volition of the former; but are not able to observe or conceive the tie, which binds together the motion and volition [...]. ([1748], Section VII, Part II)

In honor of Hume's extensive use of these attachment metaphors, I will dub them *metaphors of Humean bondage* (with due apologies to Spinoza).

These metaphors have been in continued use since the time of Hume:

> Substitute for the time honored 'chain of causation', so often introduced into discussions upon this subject, the phrase a 'rope of causation', and see what a very different aspect the question will wear. (Venn [1866], p. 320)

This content downloaded from
192.58.125.16 on Sat, 07 Jan 2023 03:46:45 UTC
All use subject to https://about.jstor.org/terms

> I propose to distinguish [...] *loose* and *tight* [causes]. A loose cause requires some third thing other than itself and its effect to bind the two together [...] a tight cause is one whose connexion with its effect is independent of such adventitious aids. (Collingwood [1940], p. 314)

> [C]ausation [...] is a very general feature [...] of the way the world works: it is not merely, as Hume says, *to us*, but also *in fact*, the cement of the universe. (Mackie [1974], p. 2)

> [C]ausal processes constitute precisely the physical connections between causes and effects that Hume sought—what he called 'the cement of the universe'. (Salmon [1984], p. 156)

> I am sure that I [...] have caused ever so many people to die [...]. Acts of mine are connected to their deaths by long chains of causal dependence. (Lewis [1986a], p. 184)

There are a number of important differences in the way that these metaphors have been used. These authors are using the metaphors to make distinct points: Collingwood is drawing a distinction; Mackie is stressing the objectivity of causation, and so on. Some authors are speaking less metaphorically than others: Salmon is claiming that causes and effects really are physically connected. Some of these metaphors are now dead: 'connection' and 'chain' have effectively acquired new technical meanings, and no longer evoke their original literal meanings; 'cement' is much more evocative (I tend to think of rubber cement or Krazy Glue® rather than the powder-based cement used in construction). Nonetheless, I believe that the widespread use of these metaphors does exhibit a certain common way of thinking about causation, one that has impeded progress in the philosophical study of causation.

2 Unpacking the metaphors

What is it about the causal relation, according to philosophers, that makes it apt for description using such metaphors? I propose to answer this question by considering what it would mean for one of these metaphors to be literally true. Suppose that causes *really were* cemented to their effects. Then, given any two events, there would be an objective fact of the matter as to whether cement is present between them. Given paradigm cases where the cement is present, we could perform chemical analyses of the cement and thus determine how the causal relation is constituted. In particular, we could learn how this cement differs from various forms of pseudo-cement. Once we had done this, we would be able to develop tests for the presence of the cement in cases where we were formerly uncertain. I think that this is an accurate parable for the way in which most philosophers think about causation. It is

This content downloaded from
192.58.125.16 on Sat, 07 Jan 2023 03:46:45 UTC
All use subject to https://about.jstor.org/terms

not, however, an accurate parable for causation itself. For definiteness, I will distill this parable into a thesis, which I dub the *Thesis of Humean Bondage*.

> **THB**: In any concrete situation, there is an objective fact of the matter as to whether two events are in fact bound by the causal relation. It is the aim of philosophical inquiry to analyze this relation.

I will argue that **THB** is mistaken, and that it has led to pseudo-problems in the theory of causation.

There are a number of ways in which one might object to **THB**. One might object to the word 'objective'. Some might say that Hume would so object, although I would not. Another 'objective' objector might be van Fraassen, who argues that it is a context- and interest-relative matter what causes what ([1980], Chapter 5). I have no quarrel with the word 'objective'. Rather, I object to the word 'the' preceding 'causal relation'. This word indicates a definite description, which presupposes that there exists a unique causal relation. The early Bertrand Russell, who gave us this understanding of definite descriptions, would object to this one because of the presupposition of existence. That is, Russell denied that there is a causal relation:

> [T]he word 'cause' is so inextricably bound up with misleading associations as to make its complete extrusion from the philosophical vocabulary desirable [. . .].
>
> All philosophers, of every school, imagine that causation is one of the fundamental axioms or postulates of science, yet, oddly enough, in advanced sciences such as gravitational astronomy, the word 'cause' never occurs [. . .]. [T]he reason why physics has ceased to look for causes is that, in fact, there are no such things. The law of causality, I believe, like much that passes muster among philosophers, is a relic of a bygone age, surviving, like the monarchy, only because it is erroneously supposed to do no harm. ([1913], p. 1)

I do not subscribe to Russell's causal nihilism. Following Cartwright ([1979]), I believe that a distinction between causal and non-causal relations grounds the essential distinction between effective and ineffective strategies. Lung cancer is correlated both with smoking and with stained teeth, but if we wish to avoid lung cancer, it will pay to quit smoking but not to whiten our teeth. My objection to **THB** lies rather with the presupposition of uniqueness.

The belief in a unique causal relation—causation—is peculiar to philosophers. For example, in Judea Pearl's landmark book on causal modeling, *Causality* (Pearl [2000]), if one looks in the index under 'causation' one will find references to numerous theories of causation mooted by philosophers, but no entry for the author's own definition. Instead, one finds definitions for a wide variety of causal concepts such as causal effect, causal relevance, total effect, direct effect, indirect effect, actual cause, contributing cause, and so on. No one of these lays sole claim to the title of causation. For

This content downloaded from
192.58.125.16 on Sat, 07 Jan 2023 03:46:45 UTC
All use subject to https://about.jstor.org/terms

Pearl and his colleagues who are in the business of constructing causal models and making causal inferences, 'causation' is a subject matter, not a relation. (Spirtes, Glymour and Scheines, whose *Causation, Prediction and Search* ([1993]) is a major contribution to the causal modeling literature, are honorary non-philosophers in this regard.)

Not all philosophers are committed to the existence of a unique causal relation. Eells ([1991]) admits different kinds of causal relations; and Lewis ([2000]) admits a certain amount of play in our ordinary conception of cause. Yet even these philosophers tend to admit varieties of causation along one dimension, while steadfastly adhering to a monolithic conception of causation along another.

The sort of pluralism about *causation* that I wish to advocate should not be confused with pluralism about *causes*, which is accepted by the majority of contemporary philosophers. The careless tossing of a lit cigarette, the recent drought, the presence of oxygen in the atmosphere; these all count among the causes of the forest fire. When talking informally, we may speak of one or another of these as 'the' cause. Which cause we single out will depend upon context and the interests of the speakers: law enforcement officials may focus on the tossed cigarette, while meteorologists focus on the drought. Nonetheless, philosophers have learned to recognize all of these factors as genuine causes. Similarly, historians may debate whether German militarism, the European alliance system, or the assassination of Archduke Ferdinand was 'the' cause of the First World War; philosophers have learned to construe these as debates over the relative importance of these causes.

3 Theories of causation

In order to bring out my objection to **THB** more clearly, it will be helpful to provide a brief overview of some philosophical approaches to causation. This will be familiar ground for most readers, yet it will be helpful to bring certain similarities and differences between these approaches into the foreground. In necessity/sufficiency approaches to causation, as well as counterfactual, probabilistic and process approaches, the analysis of causation can be broken into two distinct stages. The first stage involves the identification of some privileged class of entity, and the discrimination of the members of this class from various impostors.

In necessity/sufficiency approaches, the privileged entities are the laws of nature, and in particular, the so-called laws of 'succession'. Newton's second law of motion is an example of such a law. Laws of nature must be distinguished from accidental generalizations—e.g., 'all of the coins in Nelson Goodman's pocket are made of silver'. There are a number of proposals for making this distinction: laws support counterfactuals, are inductively

This content downloaded from
192.58.125.16 on Sat, 07 Jan 2023 03:46:45 UTC
All use subject to https://about.jstor.org/terms

supported to their instances, are formulated in terms of purely general predicates, and so on. Laws of succession are also to be distinguished from laws of co-existence such as the ideal gas law.

In counterfactual theories, the basic building blocks are counterfactual conditionals, such as this one: 'If this match had not been struck, it would not have lit.' These conditionals must be distinguished from other conditionals, such as material, strict, or indicative conditionals. Lewis ([1973b]) offers a semantics for counterfactual conditionals in terms of a metric of similarity among possible worlds. Moreover, counterfactual theories of causation employ *non-backtracking* counterfactuals, which are to be distinguished from *backtracking* counterfactuals—'If this match had not lit, it would not have been struck.' Lewis ([1979]) attempts to describe the metric of similarity appropriate for non-backtracking counterfactuals.

Probabilistic theories are formulated in terms of probabilistic correlations. Genuine probabilistic correlations must be distinguished from mere associations in frequency data. A coin that is tossed ten times with the left hand and ten times with the right may land heads on six of the left-handed tosses, but on only five of the right-handed tosses; it does not follow that the coin is more *likely* to land heads when tossed with the left hand. How to draw this distinction is a central concern in the foundations of probability. In addition, probabilistic correlations that reflect direct causal relationships must be distinguished from 'spurious' correlations that result when two factors have a common cause. The probabilistic correlation between smoking and lung cancer reflects a genuine causal influence of the former on the latter; the probabilistic correlation between stained teeth and lung cancer does not: you cannot prevent lung cancer by having your teeth whitened. Following Reichenbach ([1956]), a standard approach to this problem is to reject as spurious those correlations that disappear when we condition upon further causes of the effect in question.

In process theories, causal processes play the central role. Causal processes include physical objects like footballs and photons, as well as certain phenomena such as sound waves that are not comprised of a single object. These must be distinguished from what Kitcher ([1989]) calls 'spatiotemporal junk'—arbitrary spacetime regions that have the same 'shape' as genuine processes. Moreover, causal processes must be distinguished from 'pseudo-processes' such as shadows or spots of light. According to Salmon ([1984]), causal processes are those that are able to transmit 'marks'; according to Salmon ([1994]) and Dowe ([1992], [2000]), causal processes are those that possess conserved quantities.

This is perforce a very brief overview—readers will be familiar with many of the details. But I do wish to draw the reader's attention to the strong parallels among the problems that face the various theories in stage

This content downloaded from
192.58.125.16 on Sat, 07 Jan 2023 03:46:45 UTC
All use subject to https://about.jstor.org/terms

one. The problems of discriminating backtracking from non-backtracking counterfactuals, causal from spurious probabilities, and causal from pseudo-processes might even be thought of as different facets of the same problem. Each theory attempts to capture a notion of directed, non-accidental dependence; just those patterns of dependence that can underwrite effective strategies. There are obviously a number of important issues to be addressed here: Are existing accounts of the relevant distinctions adequate? Can these distinctions be analyzed in purely non-causal terms? What are the relationships between these distinctions? Is one of these concepts more fundamental than the others? And so on. For the purposes of this paper, I have no objection to efforts directed toward these questions.

In objecting to the presupposition of uniqueness in the definite description 'the causal relation', I am advocating a certain kind of causal pluralism. One type of pluralism would be a pluralism with respect to these various types of theory: each one has a valid perspective on the study of causation, and no one of these theories should be privileged over the others. Skyrms ([1984]) and Hitchcock ([1998]) advocate this sort of pluralism, and perhaps also Sober ([1985]).[1] While I have no real quarrel with this sort of pluralism, it is not the sort that I will advocate in the present paper. My objection here is to the assumption of uniqueness that arises in the second stage of causal analysis.

In the second stage of analysis, the causal relation is analyzed in terms of the basic units described above. For example, according to regularity theories such as Mackie ([1974], Chapter 3),[2] C is a cause of E if there are circumstances S such that relative to S and the laws of nature, C is both necessary and sufficient for E. According to probabilistic theories of the form advocated by Cartwright ([1979]), Humphreys ([1989]) and Eells ([1991]),[3] C is a cause of E if and only if C raises the probability of E relative to every relevant background condition B. The background conditions B are precisely those that need to be conditioned on in order to ensure that the probabilistic correlation between C and E is not spurious. According to Lewis's counterfactual theory of causation ([1973a]), causation is the ancestral of non-backtracking counterfactual dependence. In Lewis ([2000]), causation is the ancestral of *influence*, which involves a more fine-grained type of non-backtracking counterfactual dependence. Finally, in Salmon's process theory of causation, C is a cause of E if they are connected by a causal process or a

[1] Perhaps this form of pluralism could be captured using the metaphors of Humean bondage by stressing the *variety* of ways of binding things together. One theory of causation would correspond to cement, another to rope, yet another to chain, and so on. I will not try to develop this metaphor here.

[2] Mackie does not take this to be a complete theory of causation, but only of the regularities that provide the underpinning for causal relations.

[3] Eells offers a different account for token causation; see Eells ([1991], Chapter 6).

This content downloaded from
192.58.125.16 on Sat, 07 Jan 2023 03:46:45 UTC
All use subject to https://about.jstor.org/terms

chain of causal processes.[4] There are many variants on these analyses, but these will suffice for illustration.

Whereas I sought to draw attention to the similarities between the various theories at stage one, I wish to draw attention to the differences between the theories at stage two. The regularity theories mentioned require that causes make a difference for their effects in *some* background circumstances, while the probabilistic theories require that causes make a (specific kind of) difference for their effects in *all* background conditions. The counterfactual and process theories take causation to be the *ancestral* of the basic relation, while probabilistic and regularity theories do not. These choices are in no way forced upon the theorist by her choice of basic apparatus. A regularity theorist could impose stricter conditions upon the causal relation (as Mill [1843] does), or a probability theorist could offer a more liberal theory. A regularity or probability theorist could extend her definition to close it under the taking of ancestrals (as does Bennett [1988]), while not all counterfactual theorists do so (see e.g. McDermott [1995]).

The sort of pluralism I wish to advocate, then, is a pluralism at stage two. Once one has the basic apparatus in place—causal laws, causal correlations, non-backtracking counterfactuals, causal processes—there are a number of interesting relations that can be defined. Some of these will be important in one context, others in another. When we are asked what causes what, we may pay attention to one of these relations in one scenario, to another of these relations in a different scenario. One of these relations may be a component of one philosophically significant concept, while another is a component of another. All of these relations are causal, in a broad sense, and worthy objects of study within a theory of causation. If **THB** is true, however, more must be done. From among all of these broadly causal relations, a theory must select *the* causal relation; it must tell us conclusively, of any two events, whether they are bound as cause and effect. This further task, I will argue, is a wild goose chase.

While I have set up my argument as a critique of the second stage of extant *substantive* theories of causation, this critique will provide no succor for those who reject such substantive theories in favor of *singularist* or *primitivist* views of causation, which postulate singular or token causation as a primitive. (Tooley [1987] and Carroll [1994] defend such views.) To the extent that these theories postulate a unique causal relation (rather than primitive relations of dependence of the sort normally characterized in stage one), they are equally vulnerable to my critique.

[4] At any rate, this seems to be a necessary condition on Salmon's view: he never offers an explicit analysis of our ordinary causal locutions. Dowe ([2000]) has a more complex view that I will not consider here.

This content downloaded from
192.58.125.16 on Sat, 07 Jan 2023 03:46:45 UTC
All use subject to https://about.jstor.org/terms

My argument will rely on a series of examples, most drawn from the recent literature on causation. In each example, the 'stage-one facts' of the case are clear. That is, it is clear which non-backtracking counterfactuals are true, which events are connected by causal processes, and so on. There is no controversy over whether the cases rely on accidental generalizations, spurious correlations, backtracking counterfactuals or pseudo-processes. Yet in each case, disagreement arises as to what *causes* what. By rejecting **THB** we can cut the Gordian knot—in each case the disagreement is spurious because there is nothing to disagree about.

We must take care, of course. The following inference is not *automatically* valid: experts disagree over whether p is the case; therefore, there is no fact of the matter as to whether p is the case. There are, however, at least two reasons why the argument against **THB** ought to be accepted as persuasive. First, theories of causation are typically tested by comparing their verdicts with those of intuition. Our survey will demonstrate just how inconsistent and imprecise our intuitions are. In the face of these multiple disagreements, it becomes implausible that our intuitive causal judgments are attuned to one single objective relation. No one theory of causation can be expected to fit with all of these intuitions, because the intuitions themselves are incompatible. The cacophony of intuitions can be nicely explained, however, by postulating a variety of causal concepts and showing how different intuitions are tracking different concepts. What we should demand of a theory of causation is that it be able to identify the causal features of a given scenario that each side is attending to in making their judgments.

Second, the concept of causation is often thought to be important to philosophy because it is an ingredient in other important concepts such as explanation, prudential rationality, and moral responsibility. What I wish to suggest is that while these concepts are indeed causal in a broad sense, the causal component of these concepts can be understood in terms of stage-one facts alone; we do not need a unique causal relation in order to analyze these concepts. In each example we will discuss, there is a dispute as to whether one event or factor *really causes* another. Yet I invite the reader to ask herself whether anything of import hinges upon the answer to this question. Would our understanding of why the outcome occurred be enhanced? Would we be better placed to make decisions about how we should act if we knew? Would we be better placed to assign praise and blame? The answer to each of these questions is a resounding 'no'.

4 The two assassins

The first example is based upon one that was presented to me by Michael McDermott (personal communication) as a counterexample to the theory of

This content downloaded from
192.58.125.16 on Sat, 07 Jan 2023 03:46:45 UTC
All use subject to https://about.jstor.org/terms

Hitchcock ([2001a]). The example worried me, and I began sharing it with numerous colleagues. I found that many had no clear intuition on the central causal question presented by the case, and those that did were equally divided into the affirmative and negative camps. Moreover, those who did have fairly clear intuitions often reported that their intuitions reversed when the details of the story were presented in a different order.

The example runs as follows: Two assassins, Captain and Assistant, are on a mission to kill Victim. Upon spotting Victim, Captain yells 'fire!', and Assistant fires. Overhearing the order, Victim ducks and survives unscathed. The stage-one facts of the case are clear. For example, the following counterfactuals are all true. If Captain hadn't yelled 'fire!', then Assistant would not have fired, Victim would not have ducked, and Victim would have survived. If Victim hadn't overheard the order, or hadn't ducked, he would not have survived. These are all non-backtracking counterfactuals that accurately describe the scenario. Analogously, given the relevant background conditions, there are probabilistic correlations between yelling 'fire!' and shooting, yelling 'fire!' and ducking, and so on. These are not merely spurious correlations. The Captain's order was transmitted to both Assistant and Victim via sound waves, which are causal processes. Victim's survival was mediated by causal processes within his body that were present at the time he overheard Captain's order. And had Victim not ducked, Assistant's shot would have been connected to Victim's death by the causal process consisting of Assistant's bullet. Of these stage one facts, there is no dispute.

But let us now ask: did Captain's yelling 'fire!' cause Victim to survive? Here the agreement ends. The reader may even feel the disagreement within her own mind. If causation is indeed some type of cement that binds the events of the world together, then there must be a fact of the matter as to whether these two events do in fact stand in the causal relation. But why think there must be some such causal fact over and above the stage one facts described above? Our intuitions do not single out a correct answer to our causal question. And we do not need an answer to this causal question to answer the standard barrage of questions. Does Captain deserve praise for saving Victim's life? Clearly he does not. If Captain actually wanted Victim to survive, did he pursue a rational course of action? No more rational than the alternative of withholding the order. Do we lack anything by way of understanding of why Victim survived this incident? No we do not. What possible grounds could we have, then, for caring whether Captain's order really caused Victim's survival?

One interesting feature of this example is that it raises questions about the *transitivity* of causation. Most readers would probably grant that Captain's order caused Victim to duck, and that Victim's duck caused him to survive. If causation is transitive, as Lewis ([1973a], [2000]) maintains, then we must also

This content downloaded from
192.58.125.16 on Sat, 07 Jan 2023 03:46:45 UTC
All use subject to https://about.jstor.org/terms

grant that Captain's order caused Victim to survive. There are a number of putative counterexamples to the transitivity of causation that have a structure that is similar to the case of the two assassins. An event C introduces a 'threat' to the occurrence of E, but also initiates a process that protects E from that threat. In many such cases, we find it natural to judge that C is not a cause of E (cases of this sort can be found in Bennett [1987], Eells [1991], McDermott [1995], and Hall [2000]). Here is Eells' example: a failure in the city's power supply causes the university's backup generator to kick in, which causes the lights to be on in the campus library; but we would not judge that the power failure caused the lights to be on. What these examples have in common is that the threat to the relevant event is more 'direct' than is the process that preserves that event.[5] Until the time of the power failure, the lights in the campus library are illuminated with electricity from the city's power supply, and thus the elimination of that power supply poses an immediate threat to their illumination. By contrast, the power failure preserves the illumination of the light bulbs indirectly, by triggering the backup generator, which is programmed to generate electricity in precisely this sort of condition. The case of the two assassins is slightly different from these standard counterexamples, because the threat and the protector are equally indirect. It is this feature of the example that makes it so hard for our intuition to settle on a definite answer to the central causal question.

In fairness to David Lewis, he is less beholden to the metaphor of Humean Bondage than most. His program of Humean Supervenience ([1986b]) would seem to deny the existence of any kind of metaphysical cement that holds the events of the world together. Moreover, he has recognized a certain amount of 'play' in our concept of causation. Whether we judge C to be a cause of E may depend upon which unactualized possible alternatives to C we consider to be too remote or far-fetched to take seriously (Lewis [2000], p. 197.) This may not be an objective feature of any particular case, but rather a function of contextual factors. Yet, when it comes to the issue of transitivity, Lewis is adamant in the face of apparent counterexamples, maintaining that causation is definitely transitive. What the example of the two assassins suggests is that there is room for play in our concept of causation at this point as well.

5 The birth control pills

Our second example is from Hesslow ([1976]), who introduced it as a counterexample to probabilistic theories of causation. Thrombosis, or the forming of blood clots in the arteries, is considered to be one of the most worrisome side effects of birth control pills. This means, presumably, that the

[5] For further discussion of this point, see Hitchcock ([2001a]).

This content downloaded from
192.58.125.16 on Sat, 07 Jan 2023 03:46:45 UTC
All use subject to https://about.jstor.org/terms

consumption of oral contraceptives *causes* thrombosis. Yet among women who are fertile, sexually active, and otherwise quite capable of becoming pregnant, and who are under 35, non-smokers, and otherwise at low risk of thrombosis, birth control pills *lower* the overall probability of thrombosis. This is because birth control pills are effective preventers of pregnancy, which is itself a significant risk factor for thrombosis. Thus we have a prima facie counterexample to the thesis that causes raise the probabilities of their effects.

Do birth control pills *cause* thrombosis, or do they *prevent* thrombosis? In the case of the assassins, the reader might well have felt that neither answer to the causal question had much to recommend it; in the present example, however, the problem is an embarrassment of riches. There is a sense in which *both* answers seem correct. Once again, if causation is some kind of cement that holds the events of the world together, there must be some fact of the matter about whether or not that cement really is present in any given occasion on which a woman takes birth control pills and develops thrombosis. And once again, our ambivalence about the causal relationship between thrombosis and the consumption of oral contraceptives does not rest on any confusion about what the stage one facts of the case are. For example, the negative correlation between birth control pill consumption and thrombosis is not spurious.

There is a very natural diagnosis of why we feel compelled to say that birth control pills both cause and prevent thrombosis. The consumption of oral contraceptives affects a woman's chance of developing thrombosis along (at least) two different routes. In analogy with the concepts of net and component forces in Newtonian mechanics, we might say that birth control pills have two distinct component effects upon thrombosis (see Hitchcock [2001b]). Along one route—the one that includes pregnancy or its absence—the component effect is negative or preventative. By preventing pregnancy, birth control pills prevent thrombosis. Along the other route, the one that bypasses pregnancy, birth control pills cause thrombosis. Finally, then, the *net* effect of birth control pills on thrombosis is preventative. When we are asked whether birth control pills *cause* thrombosis, we might interpret this as a question about one or the other component effect, or about the net effect. This explains why there is no univocal answer to the question.

Some readers may be troubled that this example involves a general or type-level causal claim, rather than a claim relating particular events. This worry would be justified if the probabilities that underwrite the various causal claims reflect heterogeneity in the population of pill-users. The case assumes, however, that these probability relations obtain within some specific sub-population: women who are fertile and sexually active, non-smokers under 35, and so on. Suppose that Betty is such a woman, and suppose that Betty's probabilities for developing thrombosis are just those described in the

This content downloaded from
192.58.125.16 on Sat, 07 Jan 2023 03:46:45 UTC
All use subject to https://about.jstor.org/terms

example. In fact, Betty uses birth control pills and develops thrombosis. Did Betty develop thrombosis *because of* or *despite* her consumption of oral contraceptives? The question arises at the level of singular causation just as it did at the level of general causation, and the diagnosis is essentially the same.

Component effects and net effects are both useful causal concepts: neither is uniquely entitled to the sobriquet 'causation'. In a given context where causal information is needed, one or the other concept may be appropriate. Suppose, for example, that Carla is fertile, sexually active, and otherwise quite at risk of becoming pregnant. You are horrified to find out that she does not employ contraception. When asked why she does not use oral contraceptives, Carla replies that birth control pills cause thrombosis, and that she would prefer to avoid that dread disease. There is clearly a fallacy in Carla's reasoning: she is at greater risk of suffering thrombosis now than she would be if she were to take birth control pills. In evaluating the rationality of Carla's decision, it is the net effect of birth control pills on thrombosis that is relevant. By contrast, suppose that a pharmaceutical company that manufactures birth control pills is asked to develop a safer product. It would be an act of sophistry to rebuff this demand on the grounds that birth control pills already *prevent* thrombosis. Here it is relevant that birth control pills have a positive component effect on thrombosis; birth control pills cause thrombosis via a route that bypasses pregnancy. This raises the hope that it will be possible to eliminate this component effect of birth control pills, while preserving their effectiveness in preventing pregnancy. (Happily, this has largely been achieved since the time of Hesslow's paper.)

6 The smoker-protector gene

According to the probabilistic theory of causation advanced in Eells ([1991]), which closely follows Cartwright ([1979]), smoking causes lung cancer in a particular population just in case smoking raises the probability of lung cancer within all of the background contexts that are manifested by the members of that population.[6] Dupré ([1984]) dubs this the 'contextual unanimity' requirement, and raises an objection. Suppose that there is a rare gene which protects its bearers from the harmful effects of smoking; in fact, the handful of humans who possess this gene are slightly more likely to suffer from lung cancer if they do *not* smoke. Would we not still say that smoking causes lung cancer in the human population? Eells would be forced to deny this (although he could say that smoking is a *mixed* cause of lung cancer in

[6] More precisely, for Eells, causation is relative to a population and a population-type. Causes must raise the probability of their effects within each background condition permitted by the population-type, regardless of whether those conditions are actually instantiated in the given population.

This content downloaded from
192.58.125.16 on Sat, 07 Jan 2023 03:46:45 UTC
All use subject to https://about.jstor.org/terms

the human population, and also that it is a cause of lung cancer in the population of humans who lack the gene[7]). Dupré argues that we must replace Eells' contextual unanimity requirement with the requirement that causes raise the probabilities of their effects in a 'fair sample'—in essence, causes must raise the probabilities of their effects on average.[8] The dispute between Eells and Dupré on this issue spanned five publications (Eells and Sober [1983], Dupré [1984], Eells [1987], Dupré [1990], Eells [1991]).

Note that there is no disagreement over the stipulations of the case. Smoking increases the probability of lung cancer for those who lack the gene; it decreases the probability of lung cancer for those who have it. These are not spurious correlations, but reflect a direct influence of smoking on lung cancer.[9] And the gene in question is very rare, so that on average, smoking increases the probability of lung cancer in the human population. Eells and Dupré disagree about whether smoking *causes* lung cancer in the human population because they are employing the resources of probabilistic theories of causation to define distinct, but equally valid, causal concepts.

Dupré's concept of average effect is closely related to the concept that is called 'causal effect' in the causal modeling literature (see e.g. Rubin [1974]). This concept is very useful for a number of reasons. First, this concept corresponds to the information that is generated by controlled experiments. Ethical issues aside, suppose one were to randomly divide human subjects into two groups; members of one group would be made to smoke, while members of the other group would be prevented from smoking. The relative frequency of lung cancer within the two groups would provide evidence about the average effect of smoking on lung cancer. On the other hand, such an experiment could not tell us whether smoking is *unanimous* for lung cancer. Moreover, average effects can be used to predict the outcomes of certain kinds of interventions. For example, we might wish to know whether lung cancer rates would increase or decrease in response to a decrease in smoking rates triggered by a tax increase: they will decrease just in case smoking is a cause of lung cancer in Dupré's sense.[10] The causal concepts defined by Eells would not enable us to answer this question.

That smoking raises the probability of lung cancer on average is a contingent fact about the current human population. If it were to change in

[7] More precisely, smoking is a cause of lung cancer relative to a population-type which excludes individuals who possess the gene.

[8] Note that this notion of 'on average' probability increase does not collapse into mere correlation, as Eells ([1991], pp. 102–3) suggests. For details, see Hitchcock ([1998], [2001c])

[9] Fisher ([1959]) famously speculated that the correlation between smoking and lung cancer might be spurious.

[10] This claim is subject to a number of caveats: the tax increase is uniformly applied to all members of the population; the effect of smoking on lung cancer is not frequency dependent, and so on.

This content downloaded from
192.58.125.16 on Sat, 07 Jan 2023 03:46:45 UTC
All use subject to https://about.jstor.org/terms

such a way that the protective gene becomes more prevalent, then the average effect of smoking on lung cancer would be reversed. For Dupré, then, facts about the causal relationship between smoking and lung cancer depend upon highly contingent facts about human evolution. Eells, by contrast, is interested in defining concepts that have more the character of laws of nature. In particular, he wants a notion of causation in which causal relations are robust: they do not change in response to changes in the contingent makeup of the population in question.[11] The concept of a causal law, in contrast to causal fact, certainly seems worthy of exposition. Dupré and Eells are simply after different, but equally legitimate, causal concepts.

It should be noted that Eells is not one who is otherwise in the firm grip of **THB**. He countenances many ways in which smoking might be causally relevant to lung cancer: it might be a positive, negative, or mixed cause of lung cancer in a given population. Moreover, smoking might cause lung cancer in one population and prevent lung cancer in another (as it does in Dupré's example). Eells does not accord any privileged metaphysical status to the relation of (positive) causation: the three types of causal relevance play equal roles within his theory. Yet he was still sufficiently gripped by **THB** to worry about whether his causal concepts or Dupré's were the 'right' ones.

7 The bicycle thief

Our fourth example comes from Dretske ([1977]). Susan breaks into a sporting goods store and steals a bicycle. Later that evening, she is apprehended by the police and put under arrest. Let us grant that the following counterfactuals are true: 1) if Susan had acquired the bicycle in some other manner (if she had bought it, rented it, borrowed it, etc.), then she would not have been arrested; 2) if Susan had stolen some other item (a pair of skis, roller blades, a surf board, etc.) she would still have been arrested. I will let the reader fill in further details concerning probabilities, causal processes and so on. The question is: did Susan's stealing the bicycle cause her to be arrested?

Superficially, it would seem natural to answer in the affirmative. Note, however, that we can generate two versions of our question by stressing different parts of it: 1) did Susan's *stealing* the bicycle cause her to be arrested? 2) did Susan's stealing the *bicycle* cause her to be arrested? The answer to the first question seems to be 'yes', but intuition answers the second question in the negative. Thus our intuition can be driven either way simply by laying stress on a different part of the question.

[11] So long as the population remains consistent with the specified population-type; see note 6 above.

This content downloaded from
192.58.125.16 on Sat, 07 Jan 2023 03:46:45 UTC
All use subject to https://about.jstor.org/terms

Philosophical reaction to this sort of case has been a little different from in the previous examples. Few philosophers have attempted to argue that one or the other of these intuitions is simply mistaken; that is, the conflicting pull of our intuitive judgments has been fully acknowledged. Nonetheless, **THB** has done its mischief. If there is an objective fact about whether one event is cemented to another or not, and we acknowledge that there is a sense in which Susan's theft both did and did not cause her arrest, then it must follow that the phrase 'Susan's stealing the bicycle' corresponds to (at least) *two* different events. One of these events is cemented to Susan's arrest, while the other is not. Versions of this resolution have been advocated by Dretske ([1977]), Lewis ([1986c]), and Yablo ([1992]). One event—the one that is picked out when we lay the stress on *stealing*—is essentially a stealing and only accidentally involves a bicycle. A world in which Susan steals some other item is a world in which this event still occurs. Thus, the closest possible world in which this event does not occur is one in which she does not *steal* the bicycle, but acquires it in some legal manner. In such worlds, Susan is not arrested, and hence this event qualifies as a cause of Susan's arrest. The second event is one which is essentially an acquisition of a bicycle, but is only accidentally a stealing. This event does not cause Susan to be arrested.[12]

There are at least two problems with this standard resolution. First, it multiplies our ontology in a manner that should be avoided if at all possible. Second, our original problem returns when we ask whether the event that is essentially *a stealing of a bicycle* is a cause of Susan's arrest. Which counterfactual is relevant to the evaluation of this causal claim?

I maintain that there is nothing going on in this example that is not captured by the two non-backtracking counterfactuals stipulated at the outset. When we say that Susan's *stealing* the bicycle caused her to be arrested, we are asserting that if Susan had acquired the bicycle in some other manner, she would not have been arrested. Likewise, when we deny that Susan's stealing the *bicycle* caused her to be arrested, we withhold our assent from the counterfactual stating that Susan would not have been arrested had she stolen some other item. There is no further fact about whether Susan's stealing the bicycle caused her to be arrested. Put slightly differently, there is no absolute fact about whether Susan's stealing the bicycle caused her to be arrested. Whether it did so or not is relativized to a range of possible alternatives to Susan's stealing the bicycle, or to a specific dimension of variation. (See Hitchcock [1996a], [1996b] for further discussion.)

I conclude this review of our four central examples by noting that in many cases the intuitive tension among our causal judgments will be over-determined. Consider Hesslow's example of the birth control pills. We may

[12] Dretske ([1977]) speaks of event allomorphs instead of events with distinct essences.

This content downloaded from
192.58.125.16 on Sat, 07 Jan 2023 03:46:45 UTC
All use subject to https://about.jstor.org/terms

be conflicted about whether birth control pills cause thrombosis for a variety of reasons. As suggested in Section 5 above, it may be because birth control pills have a positive component effect on thrombosis but a negative net effect. Or it may be because birth control pills have a negative net effect on thrombosis within some subpopulations, such as healthy women who are at risk of pregnancy, while having a positive net effect on thrombosis among women who are infertile or who are heavy smokers. Finally, it may be because birth control pills have a negative net effect on thrombosis relative to an alternative in which no form of birth control is used, while having a positive net effect on thrombosis relative to other effective forms of birth control. These sources of intuitive conflict are interrelated, but not identical. It ought to be the job of a philosophical theory of causation to elucidate the nature of these interrelationships.

8 Further examples

There are a great many more examples in the literature where philosophers are in agreement about the relevant stage one facts, yet in disagreement about what causes what. Reasons of space (and the reader's limited indulgence) prevent me from discussing each in the sort of detail pursued with the four examples above. Nonetheless, a brief review of a number of cases with interestingly different structures ought to persuade the reader that the phenomenon is not isolated to a few side-show examples.

8.1 Indeterminism

A particular atom has a low probability, .01 let us say, of decaying during a certain time interval. The atom is bombarded by a neutron, increasing the probability of emission to .99. Let us stipulate that the case is genuinely indeterministic: there are no hidden variables that determine whether or not the atom will decay. In fact decay occurs. Does the bombardment cause the decay? There is no question that these are genuinely causal probabilities: the probabilistic correlation between bombardment and decay is not spurious. If the atom had not been bombarded, the probability of decay would have remained .01. The bombardment was neither necessary nor sufficient for the decay. The interaction between the neutron and the atom is a causal interaction, in the sense of Salmon ([1984]) and Dowe ([2000]). None of these facts are in dispute. So does the bombardment cause the atom to decay? Many philosophers believe that it does: the indeterministic nature of the relationship between bombardment and decay does not prevent the former from being a cause of the latter. (See especially Salmon [1984], Humphreys [1989], and Lewis [1986a].) By contrast, Kitcher ([1989]) and Hausman

This content downloaded from
192.58.125.16 on Sat, 07 Jan 2023 03:46:45 UTC
All use subject to https://about.jstor.org/terms

([1998]) deny that the bombardment is an indeterministic cause of the decay. Rather, the bombardment (deterministically) causes the probability of decay to be .99; as a matter of sheer chance, the decay does in fact occur.[13] A third possibility is that the case is underdetermined: there is some probability that the bombardment caused the decay, and some probability that the decay was spontaneous. (This position is defended by Tooley ([1987]), Woodward ([1990]), and Carroll ([1994]), and is discussed at some length in Hitchcock [forthcoming]).

8.2 Probability-lowering causes

A golf ball is rolling along the green toward a hole; it has a good chance of going in. A squirrel runs across the path of the ball, kicking it away from the hole and giving it a much lower chance of going in. This chance-lowering is not any kind of spurious correlation. As luck would have it, however, the ball comes off the squirrel's foot on a trajectory so that it bounces off a rock, back toward the hole, and in. The complete trajectory of the golf ball constitutes a causal process. Did the squirrel's kick *cause* the ball to land in the hole? Eells and Sober (Eells and Sober [1983], Sober [1985], Eells [1991]) argue that it did. (Salmon [1984] presents a similar argument using a different case.) Most people do not hesitate to say that the kick was causally relevant to, or part of the causal history leading up to, the ball's landing in the hole. But when asked whether the squirrel's kick *caused* the ball to go into the hole, many people balk. (See Mellor [1995], 67–8 for related discussion of a similar example.)

8.3 Parts vs wholes

A dry match is struck in the presence of oxygen, in the absence of a strong wind, and so on. It lights. Does the striking of the match cause it to light? Most theories of causation (e.g Lewis's counterfactual theory, Mackie's theory of *inus* conditions) would rule that it does; but there are notable exceptions. Mill ([1843]) holds that causes are unconditionally followed by their effects, hence the cause of the match's lighting is not the striking of the match *per se*, but the striking of the match, which is dry, in the presence of oxygen and the absence of a strong wind, etc. 'The real Cause,' Mill writes, 'is the whole of these antecedents' ([1843], Book 3, Chapter 5, §3). This position—or rather a probabilistic analog—has recently been endorsed by Humphreys ([1989], especially §25). No parties to this dispute disagree over

[13] Kitcher puts this forward explicitly as a position about explanation, but since he argues that causal structure is to be defined in terms of explanatory structure, it seems reasonable to attribute to him the view presented in the text.

This content downloaded from
192.58.125.16 on Sat, 07 Jan 2023 03:46:45 UTC
All use subject to https://about.jstor.org/terms

the stage one facts: the striking of the match is not, by itself, sufficient for lighting; the match would not have lit had it not been struck; nor would it have lit had oxygen not been present; and so on.

8.4 Symmetric overdetermination

Two soldiers shoot at a prisoner, their bullets simultaneously piercing his heart. Is each soldier's shot, taken individually, a cause of death? Lewis's counterfactual theory of causation has the consequence that it is not, although the mereological sum of the two events is (Lewis [1973a], [2000]). Other writers (e.g. McDermott [1995]) find this to be a failing in Lewis's account, and strive to produce accounts that yield the opposite conclusion. Once again, there is no disagreement about which counterfactuals are true. If the first soldier hadn't shot, the prisoner still would have died; likewise for the second soldier. If *neither* had shot, the prisoner would not have died. These are non-backtracking counterfactuals. Moreover, each shot is connected to the prisoner's death by means of a causal process. What is added to the case by insisting (or denying) that each shot, taken individually, is a cause of death? Pearl ([2000], Chapter 10) defines two distinct notions, 'actual cause' and 'contributing cause', such that each shot is a contributing cause but not an actual cause. Which of these two concepts is to be equated with just plain 'cause'? Philosophical readers may be surprised to learn that Pearl does not even ask this question.

8.5 Delayers

There are heavy rains in April. Since the forest is still damp in May, lightning strikes in the forest do not ignite a forest fire in May. By June, a month without rain has dried the forest and a lightning strike triggers a forest fire in June. Were it not for the heavy rains in April, there would have been a forest fire in May. If there had been a fire in May, there could not have been one in June. Did the April rains *cause* the forest fire? It is clear that if the April rains had not occurred, there would have been a fire in May and not in June. Would this have been the very same fire? That is, would *the* fire still have occurred? Bennett ([1987]) argues that an event which causes a fire to occur *later* than it otherwise would have should not count as a cause of the fire. By contrast, Paul ([1998]) argues that any event which affects the timing of an event—regardless of whether it hastens or delays the event—should count as a cause of that event. Lewis ([2000]) maintains that any event which substantially affects the time *or manner* of occurrence of another event should count as a cause of that second event.

This content downloaded from
192.58.125.16 on Sat, 07 Jan 2023 03:46:45 UTC
All use subject to https://about.jstor.org/terms

8.6 Causation by omission

The gardener neglected to water the flowers and they died. If the gardener had watered the flowers, they would have survived. Did the gardener's failure to water the flowers cause the flowers to die? Lewis ([2000]) defends such cases of causation by omission, although in earlier papers ([1973a], [1986c]) he often speaks as though causation must relate positive events. (See also Beebee [forthcoming] and Mellor ([1995], Chapter 11) for discussion.) Defenders of process theories seem committed to Lewis's earlier view: since there is no actual interaction between the gardener and the flowers, he was irrelevant to the flowers' death. (Dowe [2000] considers this to be a case of what he calls 'quasi-causation'.) Of course, the Queen of England also neglected to water the flowers: is her inaction a cause of death as well?

8.7 Double prevention/disconnection

Vandals steal a stop sign. Several hours later, an accident occurs at that intersection. The accident would not have occurred had the stop sign been there, but there was no physical process connecting the action of the vandals to the accident. Did the act of vandalism *cause* the accident? Schaffer ([2000]) argues that cases such as these are genuine cases of causation. Counterfactual and probabilistic theories of causation seem to be committed to this conclusion. Dowe ([2000]) and Hall ([2000]) have argued that the relationship between the theft and the vandalism in this sort of case is not genuine causation, but some kind of pseudo-causation (called 'quasi-causation' by Dowe and 'dependence' by Hall).

8.8 Preemptive prevention

McDermott ([1995]) describes the following example: A cricket ball is caught by a fielder; had it continued on, it would have struck a wall; had it continued on past the wall, it would have struck a window. Did the fielder's catch prevent the window from shattering? McDermott reports that while most people have the intuition that the catch did not prevent the shattering, they can be brought around by the following argument:

> If the wall had not been there, and [the fielder] had not [caught the ball], the ball would have hit the window. So between [them]—[the fielder] and the wall—[they] prevented the ball hitting the window. *Which one* [...] prevented the ball hitting the window [...]? (McDermott [1995], p. 525)

But there is no reason why the intuition that is triggered by this prompt should be viewed as correct, while a different intuition triggered by a different prompt—'because of the wall, there was no chance that the ball would shatter

This content downloaded from
192.58.125.16 on Sat, 07 Jan 2023 03:46:45 UTC
All use subject to https://about.jstor.org/terms

the window, and hence no need for the shattering to be prevented'—is viewed as incorrect. This is just another case where we have no clear judgment. As Collins ([2000]) notes, our intuitions in preemptive prevention cases can vary depending upon the nature of the auxiliary preventer. If it is a second fielder, we are more inclined to judge that the first fielder's catch prevented the window from shattering; if it is a wall, we are less inclined; and if the window is on the moon (so that the earth's gravitational field is the auxiliary preventer), we are less inclined still. Collins ([2000]) and Lewis ([2000]) conclude, on the basis of this kind of case, that our judgments of causation depend upon which non-actual possibilities we deem to be too far-fetched. It is not so far-fetched that both fielders would miss the ball, somewhat more so that the ball would avoid the wall to smash the window, and wildly far-fetched to suppose that the ball would smash a window located on the moon.

8.9 Quantitative variables

Barney Solomon had a blood pressure of x, and he suffered a stroke. If his blood pressure had been lower than x, he would have been less likely to have a stroke. If his blood pressure had been higher, he would have been more likely to have a stroke. Did his having a blood pressure of x *cause* his stroke? Eells ([1991]) argues that the answer depends upon which counterfactuals of the following form are true: if Barney's blood pressure had not been x, then it would have been x' with probability $P(x')$. Hitchcock ([1993]) criticizes this response, and argues that there is no absolute fact of the matter in this sort of case. Here is a prompting question: Does it matter what x is? If so, what is the cut-off point? That is, how do we choose an m such that: if Barney's blood pressure is below m, his having that blood pressure did not cause his stroke; but if his blood pressure is above m, his having that blood pressure did cause his stroke?

9 Conclusion

There are a great many cases where we are unclear about what causes what, even though we are clear about all the facts that are supposed to constitute causal relations. The explanation, I contend, is a false presupposition contained in the question: Do events C and E stand in *the* causal relation? There are many causal relations, and C might stand to E in some of these relations, but not in others. Here are some candidate causal relations that are brought out by our four central examples: C belongs to a causal chain of events leading up to E; C has a component effect on E along some particular causal route; C has a net effect on E when all causal routes from C to E are taken into consideration; C is a cause of E on average in some contingently

This content downloaded from
192.58.125.16 on Sat, 07 Jan 20hu, 01 Jan 1976 12:34:56 UTC
All use subject to https://about.jstor.org/terms

constituted population; C is a cause of E as a matter of causal law; C is a cause of E relative to some particular range of alternatives or domain of variation. The examples show that these relations need not be extensionally equivalent. The time has come to re-direct the resources of theories of causation toward analyzing this collection of causal concepts, and to abandon attempts to characterize *the* causal relation.

Acknowledgments

I would like to thank Charles Daniels, Alan Hájek, Laurie Paul, Jonathan Schaffer, Jim Woodward, and audience members at the Universities of Alberta, British Columbia, Washington, and Pittsburgh.

This paper was originally written for a workshop at the University of Pittsburgh in November 2000. The workshop was held in honor of Wesley Salmon on the occasion of his retirement. It was the last time I saw him prior to his tragic death in April 2001. He specifically requested a small workshop, where the papers were not to focus on his work, but on new work in his areas of expertise. At the end of the conference, he told the audience that his intention was to demonstrate that it was safe for him to retire; upon hearing the papers, however, he declared a renewed desire to continue his own work. These comments were a testament to his humility and generosity of spirit. To merely thank him for comments on this paper would be a grave injustice; as my mentor, friend, and role model, he will continue to have a profound influence on everything that I write.

Division of Humanities and Social Sciences
California Institute of Technology
Pasadena
California 91125
USA
cricky@caltech.edu

References

Beebee, H. [forthcoming]: 'Causing and Nothingness', in Collins *et al.*, [forthcoming].

Bennett, J. [1987]: 'Event Causation: The Counterfactual Analysis', *Philosophical Perspectives*, **1**, 367–96.

Bennett, J. [1988]: *Events and Their Names*, Indianapolis: Hackett.

Carroll, J. [1994]: *Laws of Nature*, Oxford: Oxford University Press.

Cartwright, N. [1979]: 'Causal Laws and Effective Strategies', *Noûs*, **13**, pp. 419–37. Reprinted in N. Cartwright (*ed.*), *How the Laws of Physics Lie* (Oxford: Clarendon Press, 1983, pp. 21–43).

Collingwood, R. G. [1940]: *An Essay on Metaphysics*, Oxford: Clarendon Press.

Collins, J. [2000]: 'Preemptive Prevention', *Journal of Philosophy*, **98**, pp. 223–34.

This content downloaded from
192.58.125.16 on Sat, 07 Jan 2023 03:46:45 UTC
All use subject to https://about.jstor.org/terms

Collins, J., Hall, N. and Paul, L. (*eds*) [forthcoming]: *Causation and Counterfactuals*, Cambridge, MA: MIT Press.

Dowe, P. [1992]: 'Wesley Salmon's Process Theory of Causality and the Conserved Quantity Theory', *Philosophy of Science*, **59**, pp. 195–216.

Dowe, P. [2000]: *Physical Causation*, Cambridge: Cambridge University Press.

Dretske, F. [1977]: 'Referring to Events', *Midwest Studies in Philosophy*, **2**, pp. 90–9.

Dupré, J. [1984]: 'Probabilistic Causality Emancipated', in P. French *et al.* (1984), pp. 169–75.

Dupré, J. [1990]: 'Probabilistic Causality: A Rejoinder to Ellery Eells', *Philosophy of Science*, **57**, pp. 690–8.

Eells, E. [1987]: 'Probabilistic Causality: Reply to John Dupré', *Philosophy of Science*, **54**, pp. 105–14.

Eells, E. [1991]: *Probabilistic Causality*, Cambridge: Cambridge University Press.

Eells, E. and Sober, E. [1983]: 'Probabilistic Causality and the Question of Transitivity', *Philosophy of Science*, **50**, pp. 35–57.

Fisher, R. A. [1959]: *Smoking: The Cancer Controversy*, Edinburgh and London: Oliver and Boyd.

French, P., Uehling, T. Jr, and Wettstein, H. (*eds*) [1984]: *Midwest Studies in Philosophy IX*, Minneapolis: University of Minnesota Press.

Hall, N. [2000]: 'Causation and the Price of Transitivity', *Journal of Philosophy*, **97**, pp. 198–222.

Hausman, D. [1998]: *Causal Asymmetries*, Cambridge: Cambridge University Press.

Hesslow, G. [1976]: 'Discussion: Two Notes on the Probabilistic Approach to Causality', *Philosophy of Science*, **43**, pp. 290–2.

Hitchcock, C. [1993]: 'A Generalized Probabilistic Theory of Causal Relevance', *Synthese*, **97**, pp. 335–64.

Hitchcock, C. [1996a]: 'Farewell to Binary Causation', *Canadian Journal of Philosophy*, **26**, pp. 267–82.

Hitchcock, C. [1996b]: 'The Role of Contrast in Causal and Explanatory Claims', *Synthese*, **107**, pp. 395–419.

Hitchcock, C. [1998]: 'Causal Knowledge: That Great Guide of Human Life', *Communication and Cognition*, **31**, pp. 271–96.

Hitchcock, C. [2001a]: 'The Intransitivity of Causation Revealed in Equations and Graphs', *Journal of Philosophy*, **98**, pp. 273–99.

Hitchcock, C. [2001b]: 'A Tale of Two Effects', *The Philosophical Review*, **110**, pp. 361–96.

Hitchcock, C. [2001c]: 'Causal Generalizations and Good Advice', *Monist*, **84**, pp. 222–46. An expanded version will be reprinted in H. Kyburg and M. Thalos (*eds*), forthcoming, *Probability is the Very Guide of Life: A Poor Folks' Guide to the Uses of Probability*, Chicago: Open Court Press.

Hitchcock, C. [forthcoming]: 'Do All and Only Causes Raise the Probabilities of Effects?', in Collins *et al.*, [forthcoming]. *Causation and Counterfactuals*, Cambridge, MA: MIT Press.]

Hume, D. [1739–40]: *A Treatise of Human Nature*. Reprinted (e.g.) in Hume [1978], pp. xi–639.

This content downloaded from
192.58.125.16 on Sat, 07 Jan 2023 03:46:45 UTC
All use subject to https://about.jstor.org/terms

Hume, D. [1740]: *An Abstract of a Book Lately Published; Entitled, A Treatise of Human Nature*. Reprinted (e.g.) in Hume [1978] pp. 641–62.

Hume, D. [1748]: *An Enquiry Concerning Human Understanding*. Reprinted (e.g.) in Hume, *An Enquiry Concerning Human Understanding*, ed. E. Steinberg, Indianapolis: Hackett, 1977, pp. xx–114.

Hume, D. [1977]: *A Treatise of Human Nature*, Second Edition, eds. L.A. Selby-Bigge and P.H. Nidditch, Oxford: Clarendon Press.

Humphreys, P. [1989]: *The Chances of Explanation*, Princeton: Princeton University Press.

Kitcher, P. [1989]: 'Explanatory Unification and the Causal Structure of the World', in Kitcher and W. Salmon (*eds*), *Scientific Explanation*, Minneapolis: University of Minnesota Press, pp. 410–506.

Lewis, D. [1973a]: 'Causation', *Journal of Philosophy*, **70**, pp. 556–67. Reprinted in Lewis ([1986d]), pp. 159–72.

Lewis, D. [1973b]: *Counterfactuals*, Cambridge, MA: Harvard University Press.

Lewis, D. [1979]: 'Counterfactual Dependence and Time's Arrow', *Noûs*, **13**, pp. 455–76. Reprinted in Lewis (1986d), pp. 32–52.

Lewis, D. [1986a]: 'Postscripts to "Causation"', in Lewis (1986d), pp. 159–213.

Lewis, D. [1986b]: 'Introduction', in Lewis (1986d) pp. ix–xvii.

Lewis, D. [1986c]: 'Events', in Lewis (1986d).

Lewis, D. [1986d]: *Philosophical Papers*, Volume II, Oxford: Oxford University Press.

Lewis, D. [2000]: 'Causation as Influence', *Journal of Philosophy*, **97**, pp. 182–97.

Mackie, J. L. [1974]: *The Cement of the Universe*, Oxford: Clarendon Press.

McDermott, M. [1995]: 'Redundant Causation', *British Journal for the Philosophy of Science*, **40**, pp. 523–44.

Mellor, D. H. [1995]: *The Facts of Causation*, London and New York: Routledge.

Mill, J. S. [1843]: *A System of Logic*, London: Parker and Son.

Paul, L. A. [1998]: 'Keeping Track of the Time: Emending the Counterfactual Analysis of Causation', *Analysis*, **58**, pp. 191–8.

Pearl, J. [2000]: *Causality: Models, Reasoning, and Inference*, Cambridge: Cambridge University Press.

Reichenbach, H. [1956]: *The Direction of Time*, Berkeley and Los Angeles: University of California Press.

Rubin, D. [1974]: 'Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies', *Journal of Educational Psychology*, **66**, pp. 688–701.

Russell, B. [1913]: 'On the Notion of Cause', *Proceedings of the Aristotelian Society*, **13**, pp. 1–26.

Salmon, W. [1984]: *Scientific Explanation and the Causal Structure of the World*, Princeton: Princeton University Press.

Salmon, W. [1994]: 'Causality Without Counterfactuals', *Philosophy of Science*, **61**, pp. 297–312.

Schaffer, J. [2000]: 'Causation by Disconnection', *Philosophy of Science*, **67**, pp. 285–300.

Skyrms, B. [1984]: 'EPR: Lessons for Metaphysics', in French *et al.* (1984), pp. 245–55.

This content downloaded from
192.58.125.16 on Sat, 07 Jan 2023 03:46:45 UTC
All use subject to https://about.jstor.org/terms

Sober, E. [1985]: 'Two Concepts of Cause', in P. Asquith and P. Kitcher (*eds*), *PSA 1984*, Vol. II, East Lansing: Philosophy of Science Association, pp. 405–24.

Spirtes, P., Glymour, C. and Scheines, R. [1993]: *Causation, Prediction, and Search*, New York: Springer-Verlag.

Tooley, M. [1987]: *Causation: A Realist Approach*, Oxford: Clarendon Press.

van Fraassen, B. C. [1980]: *The Scientific Image*, Oxford: Clarendon Press.

Venn, J. [1866]: *The Logic of Chance*, London: Macmillan.

Woodward, J. [1990]: 'Supervenience and Singular Causal Claims', in D. Knowles (*ed.*), *Explanation and its Limits*, Cambridge: Cambridge University Press, pp. 211–46.

Yablo, S. [1992]: 'Cause and Essence', *Synthese*, **93**, pp. 403–49.

This content downloaded from
192.58.125.16 on Sat, 07 Jan 2023 03:46:45 UTC
All use subject to https://about.jstor.org/terms