

Causality Reunited

Michael Strevens

To appear in *Erkenntnis*

ABSTRACT

Hall has recently argued that there are two concepts of causality, picking out two different kinds of causal relation. McGrath, and Hitchcock and Knobe, have recently argued that the facts about causality depend on what counts as a “default” or “normal” state, or even on the moral facts. In the light of these claims you might be tempted to agree with Skyrms that causal relations constitute, metaphysically speaking, an “amiable jumble”, or with Cartwright that ‘causation’, though a single word, encompasses many different kinds of things. This paper argues, drawing on the author’s recent work on explanation, that the evidence adduced in support of causal pluralism can be accommodated easily by a unified theory of causality—a theory according to which all singular causal claims concern the same fundamental causal network.

ACKNOWLEDGMENTS

For valuable comments, thanks to the anonymous referees and to Laura Franklin-Hall, who never fails to feed the fish.

1. The Disunification of Causality

1.1 *Varieties of Disunity*

Why is it so difficult to give a philosophically satisfying theory of the nature of causality? After surveying six different accounts of causation and finding each of them limited in various ways, Cartwright (2004) suggests that there is perhaps no unified theory of causality to be had because causality itself is metaphysically disunified.

Rather than try to define what it would be for causality to be disunified metaphysically, let me consider three theories of the nature—or natures!—of causal talk and of causality itself that each, in their own way, serve as paradigms of metaphysical causal disunity.

On the first, there are two or more domain-general concepts of causality, picking out two or more different kinds of relations and attaching them to the same word, ‘causation’. In many cases the criteria corresponding to these concepts agree, but in some cases they make different judgments, one saying “yes” and another “no” to the question whether this event causes that event, from which it follows that there cannot be a unified, extensionally adequate analysis of causal talk.¹ (There could be a *disunified* analysis if there is some meta-rule governing causal discourse that resolves such conflicts when they arise. A meta-rule might, for example, declare that a causal attribution is appropriate in just those cases where at least one, or sufficiently many, of the criteria apply, in which case a disjunctive extensionally adequate analysis is possible (Godfrey-Smith 2010; Psillos 2010).)

A recent proposal of this sort is Hall’s (2004) thesis that there are two concepts of causation. One concept picks out a production relation, the other

1. In some cases, as Cartwright points out, the conflict between multiple criteria may fall short of outright disagreement: one criterion may determine that a relation is causal, while the preconditions for the application of another may not obtain—so that the second, though it does not determinately disagree with the judgments of the first, fails to agree.

a difference-making relation.² Often the two go together, but sometimes they come apart. To a first approximation, Hall seems to endorse a disjunctive meta-rule for singular causal claims: we will count a causal claim of the form *c was a cause of e* as true if either of the two relations joins the relata. A relatively simple, though essentially disjunctive and thus disunified, account of singular causation is therefore in Hall's view possible.

On the second account of causal disunity, there is not even one domain-general causal relation. Rather there are many *sui generis* relations, each proper to a particular kind of causal mechanism. Cartwright, with a nod to Anscombe, suggests that such processes as feeding gasoline, sucking air, and allowing air to flow are, in a carburetor, to be distinguished in this way. They do not, then, share a single, domain-general property in virtue of which they instantiate a unified relation of causation. What they have in common that prompts us to call them causal processes is something shallow and metaphysically uninteresting—they play the same sort of role in certain forms of inference, or they have a family resemblance of some other sort. Godfrey-Smith (2010) frames a similar proposal.

The evidence for this kind of disunity is, however, rather weak. Cartwright's six inadequate accounts of causality do not partition the territory as you would expect if causality were splintered in this way: there is not one account of causation that works well for feeding, one for sucking, one for allowing flow, and so on. Rather, the boundaries are more reminiscent of Hall's production/difference-making dichotomy. Further, the cases where a given theory of causation succeeds and cases where it breaks down often involve identical processes varyingly arranged: the same account that works well for one person throwing a rock at a vase stumbles when a second person turns up and does precisely the same thing, which provides reason to think

2. Hall calls the second a dependence relation rather than a difference-making relation, but to my ear "dependence" is too broad; "difference-making" more precisely captures what Hall has in mind.

that the Balkanized causal map will divide rather than unite even the rather narrow domain of vase-directed rock-throwing. The differences between different kinds of domain-specific causal processes such as those invoked by Cartwright are surely worth investigating further, but at the present time, they seem to slice up the world in a much different way than do the competencies of the different theories of causation.

On the third account, causal discourse is freighted with pragmatic considerations and practical interests that so complicate the practice of making causal claims that no unified analysis is feasible. On this view there are not two or more distinct kinds of causation; there is just one very complex linguistic practice that is resistant to the usual methods of philosophical analysis.

The aim of the present paper is to show how, in spite of the powerful arguments that have been offered in support of the kinds of causal disunity discussed above, an account of causation as a metaphysically unified thing can be delivered.

Just as I have not tried to define what it would be in general for causality to be disunified, I will not try to define what it is for causality to be unified. I will, however, make the case that according to the particular theory I present here, causality is clearly metaphysically unified.

There are two other things I will not try to do. First, I will not attempt to convince you that my account of causation is correct. That would require attention to an enormous range of questions and counterexamples, certainly too many to include here (though a good number are addressed in Strevens (2008)). Second, I will not attempt to contest every argument that has been given for causal disunity, or to argue against every theory of disunified causality that fits the templates surveyed above—or any other template for causal pluralism. (For other varieties of causal pluralism, see the surveys in Godfrey-Smith (2010) and Psillos (2010); for some different notions of what pluralism amounts to, see Hitchcock (2007).)

My strategy is to focus on, and to attempt to undermine, the most impor-

tant positive arguments for metaphysical causal disunity, that is, arguments that give positive reasons to think that causality is by nature diverse, complex, polyvalent. I will not say much against negative arguments, that is, arguments that infer the disunity of causation from the failure of philosophical attempts so far to construct a unified account.

The two important, independent positive arguments for disunity to be considered here are Hall's argument for two concepts of causation, and a collection of arguments purporting to show that causal facts are responsive in part to facts that concern what is in the broadest sense normal, and so depend on the complex and heterogeneous domain of normative facts. This paper will sketch a theory of causation that accepts the premises of both arguments, yet that is simple, principled, and unified. Whether you think that the theory will ultimately succeed or not, then, you should conclude at least that the premises do not lead to disunity with the inevitability that some philosophers have supposed.

Both positive arguments for disunity focus on singular causation (also known as token causation or actual causation). Singular causation will therefore be my focus too; when I talk about causal claims in what follows, I mean singular causal claims of the form *c was a cause of e*. (For my story about general or type-level causation, drawing on the same ideas as the singular story, see Strevens (2008), §7.6.)

1.2 *Hall's Two Concepts of Causation*

Ned Hall begins his argument that causality has a dual, rather than a unified, nature with an enumeration of intuitively appealing theses about causation:

1. Locality: causes are connected to their effects by way of spatiotemporally continuous sequences of causal intermediates.
2. Intrinsicness: The causal structure of a process is determined by its intrinsic, non-causal character, together with the laws.

3. Transitivity: If event c is a cause of d , and d is a cause of e , then c is a cause of e .
4. Difference-making: Counterfactual dependence between wholly distinct events is sufficient for causation.
5. Omissions—failures of events to occur—can both cause and be caused. For example, Professor Schmidt's failure to feed the fish can cause their death, and be caused in turn by a blow to the head.

I have used Hall's phrasing in the list above, but what he calls the thesis of dependence I call the thesis of difference-making (see note 2). In what follows, I give my own argument in the spirit of Hall's, which is to say, I do not reproduce his reasoning exactly.

No conceivable relation, Hall argues, can fit all of these theses. Omissions are not spatiotemporally connected to the events that they cause or are caused by. Counterfactual dependence is non-transitive: it is possible for e to depend counterfactually on d and d on c , but for e not to depend on c . Finally, counterfactual dependence is determined in part by facts that are not intrinsic to (what is intuitively) a causal process. A preventer may, as a consequence, qualify as a difference-maker and so as a cause partly in virtue of facts that are not intrinsic to any process in which it participates. (For example, Professor Schmidt may be in part causally responsible for Professor Gödel's attending a department meeting in virtue of his pushing Gödel out of the path of a large boulder and thereby saving his life. Gödel's attendance depends counterfactually on Schmidt's selflessly throwing himself in the path of the rock, but this dependence and so the causal status of Schmidt's intervention depends in part on matters that are not intrinsic to the causal process running from Schmidt's push to Gödel's attendance—namely, the presence of the boulder which in actuality made no causal impression on Gödel but which came so close to transmitting its mark.) Thus there can be no unified concept of causality, because the relation picked out by this concept should satisfy all five of Hall's

theses, but no relation can do so.

The theses can be divided into two easy-to-satisfy groups, however. The first group consists of the theses (1) through (3); it can be satisfied by a relation of “production” based, perhaps, on the relation of nomological dependence, or on one of the process theories of causation (Dowe 2000). (Hall offers his own proprietary account of production.) The second group, consisting of theses (4) and (5), can be satisfied by the relation of counterfactual dependence itself. Hall’s hypothesis, then, is that we causal talkers possess at least two concepts of causation, one picking out the production relation and one picking out the relation of counterfactual dependence (or something close by).

1.3 *Causality and the Normative*

That on a given occasion practical interests play a role in determining what causes we think to mention is an old idea. Mill (1973) proposed that the distinction between causes and background conditions is of this sort: to call something a cause is to draw attention to it in the light of its local relevance; to call something a background condition is to suggest that it temporarily be ignored. The objective causal status of causes and background conditions is, on this view, exactly the same: what changes is the target of the conversational spotlight. In the aftermath of an auto accident, Collingwood (1940, chap. 31) suggested, the highway engineer may cite as the cause a tricky camber in the road while the traffic policeperson cites the driver’s reckless speed. These two experts do not disagree on the causal factors that contributed to the accident; they do not even disagree on the objective causal importance of the various factors. They merely emphasize, for professional purposes, different parts of the causal story.

A number of writers have since suggested that considerations that Mill or Collingwood would identify as practical can make a difference to whether a factor counts as a cause at all.³ Whereas for Mill, such things only determine

3. An alternative approach is to interpret the apparent influence of practical interests on

which objectively causal factors are fit to mention in a given context, on the more recent views, the context determines which factors are objectively causal.

Practical interests are sufficiently heterogeneous that, were they to play a role in determining the facts of causation, those facts would be heterogeneous. No single theory of causality could capture them, because no single theory can capture the practical interests themselves. A kind of causal pluralism would follow. (Construction of a unified theory that determines causal relations given practical interests might well be feasible, but causal reality itself would nevertheless be plural.)

Let me exemplify this possibility by looking at some recent work arguing that what is a cause of what, at least in the token sense, depends in part on prevailing norms.

The fish died, you will recall, because Professor Schmidt was unable to feed them. Proponents of the counterfactual account of causation easily make sense of the claim that Schmidt's not feeding the fish caused their death: if Schmidt had fed them, they would not have died; thus, if counterfactual dependence is sufficient for causation, Schmidt's omission is a cause. The same is true, however, for a vast array of other omissions: had anyone fed the fish, they would have lived; every one of the seven billion people on the planet, then, is causally responsible for the death of the fish, in the sense that their not feeding them caused the fish to die.

We single out Schmidt's omission as a cause. Why? McGrath (2005) argues persuasively that the criterion that distinguishes Schmidt's not feeding the fish from the multitude of other omissions makes an essential reference to norms and what is normal. In the case at hand, what matters is that Schmidt,

causal facts as due rather to some more objective factor, such as the counterfactual robustness of the events or states of affairs in question (Franklin-Hall forthcoming). Another is Schaffer's (2005) proposal that the causal relata are richer than we have supposed—they are event contrasts rather than mere events—and that the role of context is to fix the implicit aspects of these relata, rather than to determine whether between any two such relata a relation of causation holds.

unlike everyone else, has promised to feed the fish. This obligation—the fact that Schmidt is supposed to feed the fish—makes it appropriate to say that Schmidt’s omission caused the fish’s death. The lack of such an obligation on anyone else’s part prohibits us from saying the same of anyone else’s omission. Other senses of “normal” have the same power: if something characteristically happens, with or without any obligation, then its failure to happen—the failure of what normally occurs to occur—makes that something a candidate cause.

Why is this not merely a version of Mill’s move? As McGrath points out, on a Millian version of the story, though it would be appropriate to say that Schmidt caused the death and inappropriate to say that Professor Gödel (whose office is on a different floor and who dislikes fish) caused the death, it would be quite wrong to *deny* that Gödel caused the fish to die. Causal relations unilluminated by the Millian spotlight do not cease to exist; their non-existence cannot, then, reasonably be asserted. But we do deny, McGrath continues, that Gödel’s omission killed the fish. It is not merely inappropriate to say that Gödel’s omission was a cause of the fish’s death; it is factually incorrect. The normative facts thus play a role—if McGrath is right—in determining what causes what, at least as far as singular causation is concerned.

Since the normative facts that play this role are not a single kind of thing, or even a single kind of kind of thing—the “normal” embraces statistical normality, moral norms, the norms that come with proper functions in both artifacts and organisms, and natural necessity—the causal facts depend variously on the various natures of all of these things, and singular causal relations are therefore in part constituted by these things. The normal is plural, so the causal is also plural.

More recently, Hitchcock and Knobe (2009) have argued for a similar view, using as an example the following case, devised by Knobe and Fraser (2008). On the receptionist’s desk in Professor Schmidt’s department stands a coffee mug containing pens. Administrators may take pens if they need

them, but faculty are strictly forbidden to do so; they must buy their own pens. One fine morning in May, there are only two pens in the mug. An administrator takes one while Schmidt, distracted, takes the other. Not much later the phone rings. The receptionist answers, and goes to take a message, but there are no pens left. Was Schmidt causally responsible for the problem? Was the administrator?

Non-philosophers answering these questions agree that Schmidt's taking a pen caused the problem and disagree that the administrator's taking one caused the problem, even though the problem counterfactually depends on both actions. Hitchcock and Knobe explain this datum by proposing that normative facts play a role in determining singular causal relations. Among those events on which a putative effect counterfactually depends, those that deviate from the norm, or from the norms, are causes; those that conform are not.⁴ Because Schmidt was not supposed to take the pen, his taking it was a cause of the problem, in much the same way that because Schmidt was supposed to feed the fish, his not feeding them was a cause of *that* problem. If Hitchcock and Knobe are correct, then it is not only omissions or absences, but also more palpable events such as mug-to-pocket pen transfers, whose causal status depends on normative facts.

Like McGrath, Hitchcock and Knobe propose that a heterogeneous class of facts are normative in the sense relevant to causation: moral norms, statistical norms, and the norms connected to proper functioning. And like McGrath, they point out that what is determined by the normative facts is not merely which causes may appropriately be mentioned, but which counterfactual

4. Hitchcock and Knobe prefer the manipulationist version of the counterfactual account: to be a *prima facie* cause, an event must be such that the putative effect counterfactually depends on the event even when all causal factors ("variables") not on the putative causal path are held to their actual values. (Furthermore, though less important, the test for counterfactual dependence stipulates that the counterfactual non-occurrence of the event in question is due to an "intervention" in the manipulationist's proprietary sense.) Because the details make no difference to the argument, the main text ignores this significant difference between Hitchcock and Knobe's account and McGrath's account.

dependence relations are causal: their informants not only think it appropriate to identify Schmidt's action as a cause, but they deny that the administrator's act is a cause. Singular causal relations are in part determined by, hence constituted by, norms in a pluralistic sense; the causal relations themselves are therefore plural.

Why should causal facts depend on normative facts? According to Hitchcock and Knobe, the answer lies in the function of the concept of singular causation. The role of causal concepts is on their view above all to pick out factors that can be manipulated in order to alter the course of events. Something called causal structure, they hold, encompasses all those factors that could have been be manipulated to change things, whereas singular causal relations highlight factors that *should have been* manipulated to change things.

Take an event that you might wish had not occurred, then. The counterfactual dependence criterion (or manipulationist variations thereof; see note 4) picks out all factors that, if they were manipulated, would have brought about that event's non-occurrence. Some of these factors will, however, be much better targets for manipulation, in a purely instrumental sense, than others. It is better to try to change Schmidt's behavior, for example, than to try to change the administrator's behavior. More generally, it is typically best to intervene on those factors that deviate from the norm than to intervene on those that do not. (I refer you to Hitchcock and Knobe's paper for the substantial argument to this effect, including the argument that there are practical as well as moral reasons to intervene on deviations from moral norms.) The concept of singular causation is supposed to pick out only these best targets for manipulation.

If Hitchcock and Knobe are correct, it is unsurprising that the notion of singular causation is pluralistic: built into the nature of singular causal relations itself is a purely practical consideration—which manipulations are instrumentally best—on which almost any worldly fact might bear. There is an extremely general rationale for causal pluralism to be found here. According

to this rationale, causal notions have a practical function in the cognitive economy, and consequently, aspects of causal structure are entrained to the facts that influence the notions' ability to carry out this function. The more instrumental the function, the more likely such facts are to constitute a broad and heterogeneous set. In that case, what is causal will depend in part on a broad and heterogeneous array of facts, and so the causal will be, in nature, partly plural.

In Hall's kind of pluralism, there are two or more entirely distinct kinds of causal relations that can be asserted by lexically identical causal claims. On McGrath's and Hitchcock and Knobe's kind of pluralism, causality—the property in virtue of which a relation is causal—has several components, at least one of which, the property of “abnormality” or the property of “being an appropriate target for manipulation”, is plural, in the sense that it can amount to several radically different kinds of things. (Perhaps I should say in this case that causality is “multiply realizable” rather than “plural”.) On either view, there is disunity at the heart of causality: when we say that this was a cause of that, our assertion could be made true by one of several heterogeneous kinds of states of affairs. I wish to undercut both kinds of pluralism, by showing that the premises of the arguments surveyed above can easily be accommodated by a unified view of causality on which singular causal claims assert the existence of a single, homogeneous relation.

On, then, to causal reunification.

1.4 *A Theory of Singular Causal Claims*

The account of singular causal claims—claims of the form *c was a cause of e*—that I will sketch in what follows is drawn from Strevens (2008). It has three parts: an account of a relation I call *causal influence*; an account of a relation determined by the facts about causal influence that I call *causal difference-making*, and an account of a further derivative relation that I call *frameworked causal difference-making*. Everyday causal claims, you will see, are typically

claims about frameworked difference-making: to say that c was a cause of e is to say that, relative to the current explanatory framework, c made a causal difference to e .

The intuition that there are two distinct causal relations is to be interpreted as an appreciation of the difference between causal influence and causal difference-making.

The influence of norms on causal talk will be understood in terms of the influence of norms on the explanatory framework. Such norms are not a part of the nature of causality, on this view, but are rather a part of the framework; they have their effect on causal claims because such claims are implicitly relativized to the framework.

2. Causal Influence

Begin at reality's bedrock. We humans are disposed to read fundamental physics, I think, as describing a web of causal influence in which many fundamental-level facts come together to causally bring about, by way of the fundamental laws, other fundamental facts. (Perhaps we are ill advised to interpret physics in this way, but the present paper is about our actual practices of causal and explanatory attribution, not about what those practices ought to be.)

When Newton reigned, for example, we saw force as causal influence: whatever change a particle undergoes, it undergoes because it experiences a certain force. More exactly, the behavior of the particle depends on the net force on the particle together with the particle's present inertial velocity. We read this dependence as the expression of the totality of the causal influences on the particle. In a simple Newtonian world where the only force is gravitational, then, the causal influences on a massive particle at any time are the inertial velocity of the particle at that time together with the masses of all other particles in the universe.

The reach of influence can be, and is, extended by taking it to be a tran-

sitive relation: an influence on an influence on e is an influence on e .⁵ In the Newtonian world, for example, the influences on a particle include the influences on its past trajectory and on the past trajectories of the massive particles that currently affect it gravitationally; thus, the influences will span the entire past of the universe (or those parts of the universe that have mass).

The trajectories of a group of particles over a time interval are due to the sum of the influences on each of the particles over the interval; we can therefore speak of the causal production of fundamental-level events, if there is nothing to such an event over and above the movements of an assortment of individual particles within a given period of time.

In this paper I follow Hempel in calling such events *concrete events*. The concrete event of a certain vase's breaking, for example, is identical to the movements of certain particles—those particles that make up the vase—over a certain interval of time, namely, the approximate time during which the breaking occurred. A concrete event is individuated by every one of its intrinsic properties: had a shard of the vase, or even one particle in a shard, taken a slightly different trajectory, a different concrete event would have occurred.

I have characterized causal influence by Newtonian example. To complete the story, a metaphysics of causal influence is required, that is, a philosophical criterion that takes as input a fundamental theory and produces as output a specification of the connections within that theory that are relations of influence, and therefore a criterion for what counts as causal production of a concrete event.

I do not have such a metaphysics to offer—I outsource that sort of thing to the metaphysicians themselves. But I can point to three promising approaches to causal influence:

1. The process approach, as exemplified by Dowe's (2000) conserved

5. A somewhat more sophisticated notion of transitivity is needed to deal with continuous processes, including those in Newtonian physics, but this is not the place for such complications.

quantity theory.

2. The simple counterfactual approach to causation, applied solely to concrete events: a fundamental-level physical state of affairs such as a particle's having a certain mass or velocity is a causal influence on a (wholly distinct) concrete event if, had the state of affairs not obtained, the concrete event would not have occurred.
3. An approach that identifies causal influence with nomological dependence in the "right" direction. (This one requires considerable further development.)

Some further possibilities include Lewis's (2000) account of what he also calls "causal influence" and a manipulationist variant on the counterfactual approach, as described in Strevens (2008), §1.4. You should choose the approach you like best, or if you are like me, you may postpone the choice secure in the knowledge that, with many excellent options available, the notion of causal influence is on solid ground. My emphasis in this paper is not on what causal influence is, but on what can be done with it, and in particular, how it serves as the ultimate ground of singular causal claims.

Suppose we have the correct account of causal influence in hand; we can then use it to see what the influences are according to modern physics. Presumably the exchange of bosons will play a central role, as will the interaction between mass and space-time, and ultimately the nodes of influence may be values of fields at space-time points rather than particles. I will not dwell on these matters here; that I leave to philosophers such as Dowe.

However physics turns out, I expect it to give us a picture of the physical world as a densely reticulated web of causal influence relations among reality's fundamental constituents. In Newtonian gravitational theory, for example, every massive particle is influenced simultaneously by the state of every other massive particle; in electromagnetic theory, every charged particle is influenced by every other charged particle along the surface of its past light cone;

and so on. Further, every concrete event inherits the causal influences on each of its constituents. This web is the fundamental-level causal reality in which we live.

Some of our causal talk is talk of the causal influence relations that make up the web. When we say that the outer planets have a minuscule gravitational effect on events on earth, we are talking about influence, as we are when we say that events outside our past light cones have no causal connection to our lives. (Observe that we can make a claim about causal influence while knowing very little or nothing of the fundamental-level physics that grounds it—or while being quite mistaken about that physics. When a child sees a rock strike and break a vase, she can be pretty sure that the rock has a causal influence on the vase, though she has no grasp of quantum field theory. Knowing that a causal relation obtains is different from, and much easier than, knowing how or in virtue of what it obtains—and just as well, or else we would not have much, causally, to say.)

Our singular causal claims are not claims about what influences what: *c was a cause of e* does not mean that *c* was a causal influence on *e*. The conjunction of Mars and Jupiter has a causal influence on the breaking of the vase, because the aligned gravitational force of the two planets exerts a very slight pull on the pieces of the vase as they fly forth into the void. But it would be wrong to say that the conjunction of Mars and Jupiter was a cause of the vase's breaking. Such a claim implies that the conjunction's causal influence *made a difference* to whether or not the vase broke. It posits a relation, not of causal influence, but of causal difference-making, between the two events. To the nature of difference-making I now turn.

3. Causal Difference-Making

When we claim that *c* was a cause of *e*, the event *e* is almost never concrete. Rather, it is what you might call a high-level event or state of affairs. Whereas the concrete event of the vase's breaking would not have occurred had even

the smallest splinter deviated slightly from its actual trajectory, the high-level event of the vase's breaking would still have occurred if things had gone differently in any number of ways: if the shards had flown further, if they had exhibited a different pattern of spread, if the breaking had occurred a fraction of a second earlier, and so on. These variations, then, are just some of the many different ways that the very same high-level event could have occurred.

A causal influence on a high-level event will make a difference to the way that it occurs.⁶ But most such influences will make no difference to whether or not it occurs. The conjunction of Mars and Jupiter, for example, may pull the rock that breaks the vase, and the resulting shards, very slightly to the left. But with respect to the high-level event of breaking, that pull is in no way decisive: with or without the pull, the vase would have broken.

The role of causal claims is to pick out those aspects of the web of causal influence that made a difference to the occurrence of some given high-level event *e*. When *c* fits this description, we say that *c* was a cause of *e*.

Difference-makers for high-level events are typically themselves high-level, rather than concrete, events: they do not consist in a certain set of particles' having followed certain precise trajectories, but rather consist in the causal web's having certain abstract or coarse-grained properties. The ball's hitting the vase fair and square made a difference to the vase's breaking, but the precise point of impact did not. A concrete event individuated in part by the precise point of impact is therefore not a cause of the breaking; the cause is rather a high-level event individuated by the impact's being "fair and square" rather than a glancing blow or a near miss.

A theory of causal difference-making, and hence of the truth conditions for causal claims, will have two parts: a criterion for causal influence, determining what relations make up the web of influence, and a criterion for difference-making, determining which aspects of the web make a difference to a given

6. A factor influences a high-level event if it influences one or more constituents of the concrete event that realizes the high-level event (Strevens 2008, §3.23).

high-level event.

The approach to difference-making that springs most readily to mind in these times is surely the counterfactual approach. Applied to the web of influence, a simple counterfactual criterion would count an aspect c of the web as making a difference to a high-level event e just in case, had c not been present, e would not have occurred. The resulting account of causal difference-making will therefore mimic the simple counterfactual account of causal claims, on which c was a cause of e is true just in case c 's counterfactual non-occurrence would have led to e 's counterfactual non-occurrence, but with the additional requirement that c be a part of the causal web in virtue of which e is produced. The simple counterfactual criterion for difference-making consequently inherits the well-known problems of the simple counterfactual account of causal claims, most notably the mistreatment of cases of preemption, in which a backup cause would have brought about the putative effect had the actual cause not occurred.

There are a number of other promising approaches to formulating a difference-making criterion, borrowing variously from the probabilistic relevance, counterfactual, and manipulationist approaches to causation (Strevens 2008, chap. 2). I favor a test for difference-making that I call the kairetic criterion. It has a certain affinity with Mackie's (1974) theory of causal claims, with Hall's (2004) theory of causal production (one of his "two concepts" of causation), and with Baumgartner's (2008) approach to causation, except that unlike all of these it is not a reduction of causality to something else, but rather a criterion for determining which parts of fundamental-level causal reality—which parts of the web of influence—are difference-makers for a given event. My account, unlike these others, therefore posits two kinds of causal relations, one a fundamental-level relation between concrete events (and ultimately between instances of fundamental-level properties), and one a high-level relation between high-level events that is constructed from the fundamental-level relation.

To implement the kairetic criterion to find difference-makers for an event e , you first put a certain chunk of the causal web leading up to e —a certain part of the fundamental-level causal process producing e —in the form of a deductive argument with the conclusion that e occurred. (Here and in what follows, I assume that e is deterministically produced; the stochastic case is treated in Strevens (2008), §9.7.) This chunk will typically, perhaps inevitably, be less than all of causal reality: it will go back only so far into e 's past, and it will extend only so far in space from e 's locale. These limits will be marked by initial conditions and boundary conditions in the deductive argument, more on which below. The kairetic criterion will then detect difference-makers in the chunk.

The role of the deductive argument, let me emphasize, is to represent the relations of causal influence that produce e : the argument's logical structure mirrors relations of causal influence in the same way as might any other formal structure, such as a directed graph or a set of structural equations. Not every deductive argument entailing e 's occurrence represents a part of the causal production of e , of course. To do so, the premises of the argument must represent causal influences that were jointly sufficient to ensure e 's occurrence and the physical laws in virtue of which they exerted this influence. When this condition is satisfied, I say that the argument *causally entails* e , and that it constitutes a causal model for the production of e . (A more careful characterization of causal entailment is given in Strevens (2008), §3.2.)

Take a deductive fundamental-level causal model for the production of e , then. To apply the kairetic criterion for difference-making, make the model as abstract as you can without invalidating the causal entailment of e . What is left at the end of this abstraction procedure are causal difference-makers for e .⁷ These are causal influences in the part of the web represented by the

7. They are not all the difference-makers, indeed, not even all the difference-makers in the portion of the web with which you began. To find the complete set of difference-makers, you must apply the criterion to every possible chunk of the web.

model that are “minimally sufficient” to produce *e*.

For example, to find difference-makers for the vase’s breaking, construct a causal model of the events producing the breaking, which is to say, producing the concrete event that realizes the breaking. The model might take up the story a couple of seconds before the breaking, as a large and heavy rock is hurtling toward the vase. It will describe only the events in the vicinity of the rock and the vase, covering the rest of causally influential space by including a premise stating that there were no significant forces acting on the rock (or the vase) other than gravity and air resistance.⁸ But within these boundaries, the model will include every causal influence in exacting, excruciating detail. The precise position, disposition, and trajectory of every molecule, indeed every fundamental particle, in the rock will be specified. Likewise for the particles in the vase—and in the air that separates them. That is what it is to have a description of the fundamental-level causal process that produces the vase’s breaking.

Now begin the abstraction process: replace, as far as you can, premises or sets of premises with statements that are more abstract. A statement is more abstract than a set of premises if it is entailed by those premises and if its subject matter (the portion of fundamental-level reality it describes) is the same as or a subset of the subject matter of those premises. Keep abstracting until you can go no further without invalidating the causal entailment of the vase’s breaking—without either invalidating the entailment or making it non-causal.⁹ The result is a model that mentions only certain high-level, that is to say abstract, features of the rock, the vase, the air, and the laws of nature that relate them. It describes the rock, for example, not by specifying the fundamental-level properties of its fundamental-level particles, but in high-level terms, as an object of a certain approximate size, consistency, mass,

8. To use the word “significant” is to cheat a little. Better to specify precisely the net force on the rock from elsewhere without specifying its sources. The outcome will be the same.

9. One way to make the entailment non-causal is to make it merely logical—by, for example, replacing all premises with a single premise saying “The vase broke”.

and velocity. The same is true of the vase and the (minimal) resistance put up by the air between them. The model says enough to entail that the rock hits the vase fair and square, and that a strike of this sort is sufficient to shatter the vase. But it says little or nothing in addition.

What is left in the abstract model are the features of the causal process that made a difference to the breaking: the fact that the rock was fairly massive, that the air resistance was minimal, that the rock struck the vase, and so on. Any of these may be cited, then, as a cause of the breaking. The rock's striking the vase is such a cause. If we had started with a more generous portion of the causal web, including Sylvie the thrower of the rock, then Sylvie's throwing—characterized at a high level, without going into the fine details of particular muscle twitches—would also have remained at the end of the abstraction process. Sylvie's throw is also, then, a cause of the breaking. More generally, a high-level event is a cause of the breaking if there exists *some* causal model, corresponding to some section of the relevant causal web, in which that event appears when abstraction is done.

There is a constraint on abstraction that I have not yet mentioned: it must not proceed to the point of disjunction. You must not, for example, replace the description of the rock's hitting the vase with a disjunctive description specifying either that a rock hit the vase or that it was stepped on by an elephant. This would otherwise be a problem for the kairetic account: the disjunctive story meets the criteria for an abstraction, since it is entailed by the rock story and has the same subject matter. (Think of it as ascribing to the actual causal process the property of being either a rock strike or an elephant stomp; then, like the rock story, it has as its subject matter the actual causal process.) It is not often remarked that other accounts of singular causal claims have the same problem. For example, the vase's breaking counterfactually depends on the fact that it was either hit by a rock or stomped on by an elephant, on both a simple counterfactual approach and the more sophisticated approach advocated by manipulationists. The same is true for probabilistic relevance

accounts, minimal sufficiency accounts such as Mackie's, and so on.

Various solutions are possible. Lewis (1973) stipulated that the relata of causal claims must be events, and developed a notion of "naturalness" to disqualify as genuine events disjunctive states of affairs. The manipulationists relativize causal claims to a model, though they do not explain why disjunctive models are less legitimate as starting points than non-disjunctive models.

The kairetic account's solution to the problem comes in the form of an additional requirement of cohesion. A cohesive causal model is one whose fundamental-level realizers form a contiguous set in what you might think of (though see note 10) as causal similarity space. To satisfy the cohesion constraint, then, it must be possible to move from one causal process realizing the model to any other process realizing the model by a making a sequence of minimal physical changes to the starting process, without any of the changes at any point taking you outside the set of realizers. To put the requirement negatively, the realizers must not constitute what are from the fundamental-level point of view two or more distinct and separate sets in causal similarity space. The rock/elephant model fails to meet this standard: its realizers form two distinct sets, one involving stone and one involving flesh.

Much more must be said to precisify this requirement. In particular, a topology must be imposed on the set of all causal processes in order to define contiguity, or less formally, some sense must be given to the notion of a "minimal physical change" to a causal process. These issues are further discussed in Strevens (2008), §3.6 and Strevens (2012).¹⁰ It may be that my cohesion requirement is on the wrong track altogether, and that I would be better off borrowing Lewis's notion of naturalness. To some extent, I think, the disjunction problem can be solved independently of the development of

10. Despite my characterization—for expository purposes—of the relevant space as a similarity space in the previous paragraph, a topological structure is sufficient to define contiguity and therefore cohesion; thus, a similarity metric is not required. There need not be a fact of the matter, then, as to the degree of similarity between two different causal processes. A notion of a "minimal physical change" is enough.

the rest of the kairetic machinery. But I will say no more here, as my response to causal pluralism does not depend on the details.

To apply the kairetic criterion as it is laid out above would be a formidable undertaking. You would require a comprehensive understanding of fundamental physics, a complete knowledge of the relevant initial and boundary conditions, and considerable computational resources. No one among us could, I think, come close to executing the steps I have described. Yet even physics ingénues and mathematical novices make competent judgments about causality. Everyone can see that the rock's hitting the vase was a cause of its breaking. How can that be?

The answer, of course, is that the ultimate criterion for difference-making and the procedures we humans use to determine difference-making need not coincide. The kairetic criterion tells you what counts as the correct answer to the question "What are the causes of e ?". How you find that answer is up to you. There are many heuristics for determining difference-makers that do not require a comprehensive knowledge of fundamental physics. Some of these heuristics are built into the folk physics, or if you like the physical intuition, that comes to us either as an innate endowment or as a very early acquisition (Baillargeon 2004). They are extremely reliable, but like all heuristics they are also in principle fallible. If you use the heuristics alone you are never certain that the rock, rather than an ultrasonic pulse that arrived milliseconds before the rock, broke the vase; to determine the answer for sure, you might have to make extremely precise physical measurements, as the kairetic criterion suggests.

Is it strange that we should operate with a notion of difference-making that implies such a difficult-to-implement criterion? It should not be: it happens everywhere. Our concept of a dog, for example, apparently locates doghood in some abstract property of animal genetics, or of historical lineage, or of some combination of the two, or of something more complex still. But we very reliably judge that animals are and are not dogs without examining either

their evolutionary pedigree or their DNA. Apparently, our minds take a certain pleasure in esoteric and epistemically distant criteria for category membership, while at the same time getting by from day to day using entirely exoteric, quite practical, and very reliable rules of thumb.

Can omissions or absences be causal difference-makers, and therefore causes? Yes: after the kairetic account's abstraction procedure has been applied, a piece of the causal web may be ascribed negative properties, and these may properly be cited as causes. In the case of the rock and the vase, for example, the abstract model will specify that there were no significant forces on the rock other than the earth's gravity. The absence of such forces—forces sufficient to deflect the rock from its vase-ward trajectory—is therefore one of the causes of the vase's breaking. Whether this implies that, like Professor Schmidt's failure to feed the fish, the similar failures of billions of other potential feeders are causally responsible for the fish's dying, is a topic for the next section.



Let me pause now to respond to Hall's argument for causal pluralism. That argument, you will recall, proceeds by showing that certain intuitively compelling theses about causation cannot be satisfied simultaneously by a single relation; rather, the theses fall into two groups, corresponding to conceptually and (presumably) metaphysically distinct relations.

I believe that the argument is correct in every important way. There are two kinds of causal relation. One has the properties of locality, intrinsicness, and transitivity, while the other is a difference-making relation that allows omissions and other absences to play a causal role, and for the existence of which counterfactual dependence is typically sufficient. The first is causal influence; the second is causal difference-making.^{11,12} But this causal duality

11. On the sufficiency of counterfactual dependence for causal difference-making, see Strevens (2008) §3.81.

12. Kairetic causal difference-making is, unlike counterfactual dependence, transitive. The thesis of transitivity is therefore true of both varieties of causal relation. Problem: the relation asserted by causal claims is widely thought not to be transitive. The resolution of this

does not amount to causal pluralism. Why not?

On my picture, the world has a single, objective, homogeneous causal structure: the web of influence. (The homogeneity consists in each strand of the web having the same essential nature, to be specified by the correct metaphysics of influence: conserved quantity transmission, nomological dependence, or whatever.) Where does causal difference-making fit in, then? That one high-level event makes a difference to another high-level event is not a fact that exists over and above the web of influence, constituting a distinct kind of causal being; it is, rather, an abstract fact about the web of influence itself. Facts about difference-making stand to facts about influence in the same way that facts about centers of mass stand to facts about mass: just as a fact about a center of mass summarizes certain physically pertinent information about the underlying distribution of mass, so a fact about difference-making summarizes pertinent information about the fundamental stuff of causation, the network of relations of influence.

Why do we need such a summary? Facts about influence are, though fundamental, so unwieldy as to be useless to beings such as ourselves. We care about high-level events—birth, death, falling in love, publication—but not about their precise realization, at least, not with respect to the fundamental level, where realizers are distinguished by the minor movements of peripheral particles. Consequently, there is little to be gained by our enumerating the many fundamental-level facts that brought about that precise realization. What interests us are the aspects of fundamental-level causal reality on which the high-level events hinge, that is, the aspects that make a difference to whether or not the high-level events occur. Our difference-making criterion is constructed to sift these aspects from the rest, and our singular causal claims to advertise them far and wide.

My two causal relations, then, do not imply as Hall's do distinct, coexisting causal structures. There is, at base, a single kind of causal stuff—influence—

problem is given in Strevens (2008), §6.5.

with a single kind of fundamental structure—the web. There are, however, two ways of describing the web. The first is by stating facts about influence; these may be used to provide a complete inventory of causal reality, without remainder. The second is by stating facts about difference-making. Such facts do not add anything to what would be contained in a complete enumeration of the connections in the web, any more than facts about center of mass add anything over and above what is contained in a complete description of the universe's distribution of mass. The distribution of mass fixes all facts about centers of mass; likewise, the web of influence fixes all facts about causal difference-making. In this sense, the web of influence is all, causally, that there is, just as the mass distribution is all, massively, that there is. But the pertinent summaries of certain aspects of causal reality given by difference-making facts are extremely useful to us in the same way that facts about centers of mass can be extremely useful to physicists. Their function is not to add something—there is nothing to add—but rather to subtract something, namely, reams of detail that have no relevance to the predictive, manipulative, or explanatory purpose at hand.

In short, claims about difference-making, which means all or almost all claims of the form *c was a cause of e*, do not point to a separate causal reality existing alongside the web of influence, but rather pick out those properties of the web—the one and only causal reality—that are systematically most important to us.

4. Frameworked Difference-Making

Singular causal claims, I have said, are assertions of causal difference-making. Now let me take that back. Causal claims are normally assertions of *frameworked* difference-making: *c was a cause of e* normally says that *c* was a causal difference-maker for *e* relative to an implicitly specified explanatory framework.

The truth of a frameworked causal claim depends on two things: the facts

about causal difference-making and the facts about the framework. If the truth of a frameworked claim depends on the moral facts, it does not follow that causal difference-making depends on moral facts: it might be, rather, that the framework depends on moral facts. That, I will propose, is exactly what is going on in the cases that McGrath and Hitchcock and Knobe present to argue that causation itself is normatively freighted.

What is it to make a framework-relative claim of difference-making? In answering this question, I can be brief, because my notion of a framework is in most respects identical to Mackie's (1974) notion of a causal field. A framework is a set of propositions about the web of causal influence. Such propositions are presumed to be true by the issuers and recipients of claims of frameworked difference-making, and in what follows I will assume that they are true. (When they are false, a presupposition of the corresponding claim fails.)

To determine the difference-makers for an event e relative to a given framework, you proceed as follows. As before, take a portion of the causal web in which e is embedded, that is, a portion of the fundamental-level causal process that generated e . There is now, however, a constraint on your choice: the part of the web chosen must include the elements that make the framework's propositions true. The propositions in the corresponding causal model must therefore represent the framework's truthmakers (and hence must entail the propositions in the framework).

Now implement the kairetic procedure for determining difference-makers, that is, make the causal model as abstract as you can without either invalidating the causal entailment of e or violating the requirement of cohesion. There is one further constraint on abstraction: you must not remove the representation of the framework's truthmakers. Whatever details are removed from the model in the course of abstraction, then, the details specified by the framework must remain. Once abstraction is finished, anything left in the model is a causal difference-maker relative to the framework *unless* it is part of the framework

itself.

The effect of frameworking, then, is as follows: an event's difference-making power is always evaluated relative to a fixed background specified by the framework, but the elements of the background do not themselves count as difference-makers (though they may well be difference-makers relative to another framework or in the unrelativized sense).

Suppose, to take the usual example, that you ask for the causes of a house fire relative to a framework that specifies the presence of oxygen. You hold the presence of oxygen fixed when abstracting to determine difference-making, so that oxygen will be present in any relevant causal model, but you are prohibited from counting oxygen itself as a difference-maker. Relative to an oxygenated framework, the short circuit is a causal difference-maker for the fire but the presence of oxygen is not. Thus it is true that the short circuit was a cause of the fire, but false that the presence of oxygen was a cause of the fire, relative to the framework.

Observe that frameworks do not merely determine what it is appropriate to say, but what it is true to say. Relative to the oxygenated framework, it is not only inappropriate to say that oxygen was a cause of the fire; it is false. If our causal discourse is largely implicitly framework-relative, then, the singular causal claims we make and deny will reflect not only the causal structure of the world but also whatever facts and other considerations determine the frameworks with respect to which we talk.¹³

Now to apply the notion of frameworked difference-making to the cases claimed by McGrath and Hitchcock and Knobe to show that causality is inherently normative.

Both Professor Schmidt and the administrator took pens from the receptionist's desk; no pen was left when the receptionist went to take an important

13. Perhaps it is better to regard causal claims that cite frameworked factors not as false but as semantically defective in some other way, like claims with false presuppositions. This is an important issue for the frameworker, but I do not have the space to treat it here.

message. Knobe and Fraser's informants generally agreed that Schmidt's taking a pen was a cause, and disagreed that the administrator's taking a pen was a cause, of the problem. How can that be? The only symmetry-breaker is normative: the administrator was permitted to partake freely of the departmental writing equipment, whereas Schmidt was not. So the truth of the causal claims in question must depend on normative facts. Hence, the relation of singular causation depends on the normative facts.

I agree with this diagnosis up to, but not including, the final sentence. Here is why. The relevant framework for a causal claim typically includes, I propose, what you might call the causal background, by which I mean all the causally relevant facts that normally hold true and that in addition held true in the particular case in question. The *normally* here has just as wide a scope as McGrath and Hitchcock and Knobe imply: it includes statistical norms, functional norms, moral norms, and perhaps more. For pen-related causal claims in Schmidt's department, then, it includes in particular the morally sanctioned and statistically normal taking of pens by administrators. The causal background for any particular pen-related claim in the department will therefore contain any actual pen-takings by administrators, but not pen-takings by the faculty. This background is placed in the framework, and relative to such a framework, Schmidt's pen-taking but not the administrator's counts as a cause.

The same treatment accounts for the role of norms in determining which claims about the causal power of omissions may truthfully be asserted. As explained in section 3, an omission can be a causal difference-maker. In the case of the starving fish, the unframeworked difference-maker is not any particular omission but rather a more general fact: no one fed the fish. Thus, a cause of the fish's death was no one's feeding them.

Schmidt failed to feed the fish; so did Professor Gödel. But Gödel was not expected to feed the fish, so his not feeding them is a part of the causal background. The same goes for everyone else except Schmidt. The expected non-

feedings—the “normal” non-feedings, in McGrath’s sense of the word—are, if they obtain, a part of the causal background, and so go into the explanatory framework with respect to which causal claims are by default made. With all of these non-feedings in the framework, it is no longer permissible to cite no one’s feeding the fish as a cause, because this event or state of affairs has among its truthmakers frameworked events such as Gödel’s non-feeding. A narrower event must be cited, one that captures that part of the content of “no one fed the fish” that remains once the frameworked events are subtracted. That event is, of course, Schmidt’s not feeding the fish. Relative to the relevant framework, then, Schmidt’s not feeding the fish is a cause of the fish’s expiration, but Gödel’s not feeding them is not. The difference in framework-relative causal status is, as McGrath claims, indeed due to normative facts. Such facts have their effect, however, not because they bend causal reality but because they permeate the framework.

In my treatment of these cases, I have relied on the assumption that whatever is normal (in McGrath’s and/or Hitchcock and Knobe’s sense) is placed by default into the framework with respect to which causal claims are evaluated. There are many other ways that events can enter the framework. In prevention talk, for example, a contextually salient threat is placed into the framework: when I say that Schmidt’s push prevented Gödel’s death, I evaluate his action’s causal status against a background in which an oncoming boulder is held fixed. Some attributions contain implicit presuppositions and thus frameworkings: when I talk of the causes of a person’s “survival”, I imply and therefore framework a threat to their continued existence. Thus when I say that Schmidt’s well-timed shove caused Gödel’s survival, I framework the boulder careening down the trail. More generally, any device that creates a conversational presupposition is apt to put that presupposition, if it concerns causal structure, into the framework.

Elements may also be removed from the framework, as when a presupposition is canceled: “The short circuit caused the fire.” “What about the presence

of oxygen?” “Yes, that too was a cause.” Here, asking about the presence of oxygen puts oxygen into play as a possible cause and so removes it from the framework, thus securing its framework-relative causal status.

A complete treatment of the framework is not something I will try to give here—I believe that such a treatment is better given by linguists and philosophers of language than by metaphysicians and philosophers of science. Suffice it to say that the proliferation of human interests motivates an array of frameworks as splendid and various as anything a pluralist might desire.

The malleability of explanatory frameworks suggests a test for my view of the relation between normative and singular causal facts: changing the framework without changing the underlying normative facts should, if I am correct, change the truth values of certain causal claims. Schmidt’s, but not Gödel’s, non-feeding of the fish caused their death. Can you change the framework so that Gödel’s, but not Schmidt’s, non-feeding caused the fish’s death? That is very difficult, because a framework that assumes non-feeding by everyone except an apparently randomly chosen professor with no relation to the fish is so unusual a thing. What you can do is remove all non-feedings from the framework: “The fish died because Schmidt did not feed them.” “You single out Schmidt, but no one whatsoever fed them—correct?” “That’s right, the fish died because no one fed them.”¹⁴ The possibility of the latter causal attribution is, I think, extremely difficult for McGrath (or for Hitchcock and Knobe) to explain. The normative facts remain the same—it was Schmidt, and Schmidt alone, who was supposed to look after the fish—yet a little conversational massage implicates us all. This suggests that the relevance of norms to causes is not direct; it rather goes by way of a middle term that may or may not point to norms depending on the context (though it does so by default). That middle term is the explanatory framework.



14. The conversation might continue: “So they died because Gödel in particular did not feed them?” “Not because of Gödel in particular; because of everyone’s failure to feed them.”

The truth of a frameworked singular causal claim depends, according to my view, on two sorts of fact: a non-causal fact concerning the contents of the framework, and a causal fact concerning difference-making (and hence ultimately concerning the structure of the web of causal influence). The dependence of causal claims on normative facts arises, I have argued, from the dependence on normative facts of the contents of the relevant explanatory framework. Causal reality itself is in no way normative; the heterogeneity of norms does not, therefore, metaphysically disunify causation.

More generally, that the truth of causal claims depends on context does not imply that causal relations are inherently contextual, any more than the dependence of the truth of “I broke the vase” on the identity of the utterer shows that what causes what depends on who is speaking.¹⁵

To a certain extent, Hitchcock and Knobe might agree with this diagnosis: they identify a feature of the world, causal structure, that is determined entirely by non-normative facts. This non-normative causal stuff is distinct from singular causation, however: what they mean by the causal structure of a scenario is something like the scenario’s instantiation of a set of type-level causal relations; in cases of preemption, for example, both actual and backup causes have equal status within the causal structure. On Hitchcock and Knobe’s view, singular causal relations, and so the causal facts that distinguish actual causes from backup causes, are irrevocably shot through with normativity.

My view, by contrast, identifies an objective, unified, non-normative singular causal structure to the world. To see the difference, consider two scenarios. One is the usual case of late preemption: first Sylvie then Bruno throw their rocks at the vase; Sylvie’s rock breaks the vase while Bruno’s slightly later throw merely skims the scattered shards. The other is a version of Hitchcock and Knobe’s pen-taking case. Administrators regularly send complex copy

15. A reader asks: if a fact helps to determine the truth of a causal claim, is it not thereby causal? Answer: no. Consider the causal claim “The vase was broken by the rock if Goldbach’s conjecture is true, or by the elephant otherwise”. The truth of the claim hinges in part on the facts about Goldbach’s conjecture, but these are in no interesting sense causal facts.

jobs to the departmental printer downstairs. The printer can deal with these jobs provided that it does not have to deal at the same time with other jobs containing involved mathematical typesetting. For this reason, faculty are supposed to send their logic papers to the upstairs printer only. One afternoon, an administrator sends a complex job to the printer. At the same time, the addled Schmidt accidentally sends his proof of a new incompleteness theorem to the downstairs, rather than the upstairs, printer. The printer is overloaded and malfunctions, failing to run the administrator's job. I take it that, as with the pens, we attribute causal responsibility for the printer's breaking to Schmidt but not to the administrator.

If Hitchcock and Knobe are correct, then Bruno's throw and the administrator's printing job are non-causes of their respective breakings in exactly the same sense. They attain this status in different ways—on the one hand, through the interventionist machinery for ruling out backup causes (note 4), and on the other hand, through the normative machinery for ruling out causes that it would be inappropriate to manipulate—but the end result is identical: as potential singular causes, they are declared utterly ineffectual.

To causally equate the two seems wrong to me. The printing job tangles with and trips up the printer, whereas Bruno's rock goes nowhere near the (intact) vase. To speak in more philosophical terms, there is a structure to the causal history of the world—a structure at the level of singular, not type, causation—in which the administrator's printing job is causally connected to the printer's breaking in a potent, productive way that Bruno's throw is not connected to the vase's breaking. Our singular causal claims are claims about this difference-making structure. If they fail to mention the printing job, it is not because the job is not a causal difference-maker, but for some other reason.

That reason is, as Mill proposed, purely practical: our singular causal claims are for conversational or other contextual reasons made relative to a framework that rules certain elements in the difference-making structure

provisionally ineligible as targets of the claims. This view captures, I think, much better than Hitchcock and Knobe's, the important difference between the two cases.

5. Causality Unified

There are two kinds of arguments for causal pluralism: positive arguments, and negative arguments that turn on despair in finding a single relation that captures the intricate truth conditions for causal claims. I am not at all desperate, because I think my view can handle the truth conditions rather well (Strevens 2008, chap. 6). But powerful positive considerations in favor of plurality demand a response.

The arguments proposed by Hall and by McGrath and Hitchcock and Knobe are powerful because their major premises are true. There are in our causal conception of the world two different kinds of causal relations, and the truth of causal claims does depend in part on normative facts. But as I have shown, causal pluralism does not follow. The two causal relations of influence and difference-making do not cross-classify causal reality, but rather correspond to concrete and abstract descriptions, respectively, of a single reality, the web of causal influence. Causal claims may depend on normative facts, but causal reality does not: the component of a causal claim that exposes it to the normative realm is not itself causal.

I conclude that causality is a single, homogeneous stuff, independent of our interests. That does not mean that it is independent of our minds—everything I have said above is consistent with a Kantian outlook on the causal—but it does mean that the causal structure of the world, at both the type and token levels, constitutes a fixed and unified backdrop against which we can stage our show.

References

- Baillargeon, R. (2004). Infants' physical world. *Current Directions in Psychological Science* 13:89–94.
- Baumgartner, M. (2008). Regularity theories reassessed. *Philosophia* 36:327–354.
- Cartwright, N. (2004). Causation: One word, many things. *Philosophy of Science* 71:805–819.
- Collingwood, R. G. (1940). *An Essay In Metaphysics*. Oxford University Press, Oxford.
- Dowe, P. (2000). *Physical Causation*. Cambridge University Press, Cambridge.
- Franklin-Hall, L. R. (Forthcoming). Explaining causal selection with explanatory causal economy: Biology and beyond. In C. Malaterre and P.-A. Braillard (eds.), *How Does Biology Explain?* Springer-Verlag, Heidelberg.
- Godfrey-Smith, P. (2010). Causal pluralism. In H. Beebe, C. Hitchcock, and P. Menzies (eds.), *Oxford Handbook of Causation*. Oxford University Press, Oxford.
- Hall, N. (2004). Two concepts of causation. In J. Collins, N. Hall, and L. A. Paul (eds.), *Causation and Counterfactuals*. MIT Press, Cambridge, MA.
- Hitchcock, C. (2007). How to be a causal pluralist. In P. Machamer and G. Walters (eds.), *Thinking about Causes: From Greek Philosophy to Modern Physics*. University of Pittsburgh Press, Pittsburgh.
- Hitchcock, C. and J. Knobe. (2009). Cause and norm. *Journal of Philosophy* 106:587–612.

- Knobe, J. and B. Fraser. (2008). Causal judgment and moral judgment: Two experiments. In W. Sinnott-Armstrong (ed.), *Moral Psychology*, volume 2, pp. 441–448. MIT Press, Cambridge, MA.
- Lewis, D. (1973). Causation. *Journal of Philosophy* 70:556–567. Reprinted in *Philosophical Papers*, volume 2. Oxford University Press, Oxford.
- . (2000). Causation as influence. *Journal of Philosophy* 97:182–197.
- Mackie, J. (1974). *The Cement of the Universe*. Oxford University Press, Oxford.
- McGrath, S. (2005). Causation by omission: A dilemma. *Philosophical Studies* 123:125–148.
- Mill, J. S. (1973). *A System of Logic*. 8th edition. Routledge, London.
- Psillos, S. (2010). Causal pluralism. In R. Vanderbeeken and B. D’Hooghe (eds.), *Worldviews, Science and Us: Studies of Analytical Metaphysics: A Selection of Topics From a Methodological Perspective*, pp. 131–151. World Scientific, Singapore.
- Schaffer, J. (2005). Contrastive causation. *Philosophical Review* 114:327–358.
- Strevens, M. (2008). *Depth: An Account of Scientific Explanation*. Harvard University Press, Cambridge, MA.
- . (2012). Précis of ‘Depth’ and reply to critics. *Philosophy and Phenomenological Research* 84:447–460, 492–505.