

Sensitive and Insensitive Causation

James Woodward

California Institute of Technology

1. Dependence and Sensitivity

Consider a paradigmatic causal transaction. Suzy stands in front of a fragile glass bottle with a large rock in her hand. No other possible causes of the bottle's breaking—no backup or preemptive throwers, no earthquakes and so on—are waiting in the wings. Suzy throws; the rock strikes the bottle squarely, and it shatters. The impact of the rock caused the bottle to shatter.

In discussing such examples, philosophers who are sympathetic to counterfactual theories of causation have tended to focus on the *counterfactual dependence* of the effect on its cause. The simplest such treatment, which forms the basis for David Lewis's early (1986) account of causation, takes as its point of departure the counterfactual dependence of the *occurrence* of the effect on the *occurrence* of its cause. Thus, in the case under discussion, whether or not the shattering occurs counterfactually depends on Suzy's throw, and one might take this fact to underlie our judgment that the throw causes the shattering. More recently, Lewis (2000) has drawn our attention to other patterns of counterfactual dependence having to do with what he calls "influence"—these

Earlier versions of this essay were delivered as talks at a conference on explanation in the natural and social sciences held in Ghent, Belgium, in 2002, at Stanford University, and at the 2003 Logic, Methodology, and Philosophy of Science meetings in Oviedo, Spain. I am grateful to the audiences at these locations and especially to Johan von Benthem, Peter Godfrey-Smith, and Elliott Sober for very helpful comments and discussion. Special thanks are also due to Chris Hitchcock for detailed comments on an earlier draft and to the anonymous referee for the *Philosophical Review*.

Philosophical Review, Vol. 115, No. 1, 2006

DOI 10.1215/00318108-2006-001

© 2006 by Cornell University

describe the counterfactual dependence not just of the occurrence of the effect on the occurrence of the cause, but also the counterfactual dependence of the time and manner of occurrence of the effect on the time and manner of the occurrence of the cause. Thus, as Lewis notes, in the example just described, it ordinarily will be true that if Suzy were to throw the rock in a sufficiently similar way, but slightly earlier or later, the shattering would occur slightly earlier or later; if Suzy were to vary the direction or momentum of the throw, but in such a way that the rock still struck the bottle sufficiently hard, corresponding variations in the effect (in the manner of shattering, the dispersal of the glass, and so on) would result.

The counterfactual dependence of effects on their causes is such an obvious feature of many examples of causation that it is easy to miss the fact that there is another feature having to do with counterfactual structure that plays an important role in such examples. This feature has to do with the *sensitivity* of the causal relationship (and, more specifically, the sensitivity of certain of the counterfactuals associated with it) to changes in various other factors. Broadly speaking, a causal claim is sensitive if it holds in the actual circumstances but would not continue to hold in circumstances that depart in various ways from the actual circumstances. A causal claim is insensitive to the extent to which it would continue to hold under various sorts of changes in the actual circumstances. The sensitivity of counterfactuals is understood similarly.

I take the basic idea of sensitivity from Lewis (1986), although my detailed understanding of the notion and the uses to which I will put it will be somewhat different from his. I will argue that just as we expect that effects should depend counterfactually on their causes, at least in paradigmatic cases, so also we expect that causal relationships, at least in paradigmatic cases, should not be too sensitive. Sensitivity is a relatively neglected dimension of causal claims, but it importantly influences our judgments: we tend to regard causal claims that are highly sensitive as defective, nonstandard, or at least importantly different from less sensitive causal claims. I intend this as an empirical claim about our actual practices of causal judgment in ordinary life, but, as we shall see, it also fits with certain normative judgments we make about causal claims in the special sciences.

For the purposes of this essay, I will stipulate, following Lewis, that the occurrence of e is counterfactually dependent on the occurrence of c if and only if the following two counterfactuals hold:

Sensitive and Insensitive Causation

$$(1.1) \quad O(c) \rightarrow O(e)$$

$$(1.2) \quad \neg O(c) \rightarrow \neg O(e),$$

where $O(c)$ and $O(e)$ are the propositions that the events c and e occur. Lewis (1986) holds that counterfactual dependence is sufficient but not necessary for causation; he identifies causation with the ancestral of counterfactual dependence. I will not assume that causal claims and counterfactual dependence are always connected in the specific way that Lewis supposes, but I will assume that there are systematic connections between causal claims and claims of counterfactual dependence. In particular, I will assume that in simple cases in which the possibilities of overdetermination, back-up causes that will act if the actual cause does not, and so on are excluded, the sensitivity of the *causal* claim that c caused e will depend on the sensitivity of the *counterfactuals* (1.1) and (1.2).

As intimated above, I take the sensitivity of the counterfactual (1.1) to have to do with whether it (or, alternatively, an analogue to it relating c -like events to e -like events) would continue to hold under various sorts of changes or departures from the actual state of affairs which are such that c or a c -like event continues to occur. The reference to whether analogues to (1.1) would continue to hold is meant to allow for the possibility that some departures from actuality (for example, those that alter the time or place of occurrence of c and e) may be such that they would affect the identities of c and e . I assume that in such cases we can still ask whether the following counterfactual claim, which I take to be the analogue to (1.1), is true:

$$(1.3) \quad \text{If a } c\text{-like event had occurred in those circumstances, an } e\text{-like event would have occurred.}$$

To simplify the exposition, I will usually consider, in assessing the sensitivity of the counterfactual (1.1), whether (1.1) itself continues to hold under various sorts of changes, but the reader should think of my discussion as including the possibility that the counterfactual to be evaluated for sensitivity may be of form (1.3) rather than of form (1.1).

In assessing the sensitivity of a counterfactual like (1.1), the guiding idea is that we are to consider neighborhoods of possible changes from the actual world in which c or a c -like event occurs and then ask whether (1.1) or its analogue (1.3) would be true in such neighborhoods. To the extent that (1.1) or its analogue would continue to hold in all or many such neighborhoods (or at least in those that do not represent too

“large or substantial a departure” from actuality or in those that are most “relevant” or “important”—see section 3 below for more on all of these notions), the more insensitive the counterfactual relationship (1.1) will be; the “smaller” and so forth the neighborhoods in which (1.1) or its analogue holds, the more sensitive the relationship.

Given Lewis’s semantics for counterfactuals, (1.1) will be automatically true as long as c and e occur—an assumption that I also will adopt for the purposes of this article.¹ Thus, in such a case, whether e is counterfactually dependent on c will depend just on the truth value of (1.2). But while the counterfactual (1.1) is trivially true if c and e occur, it is a further and nontrivial question whether and to what extent (1.1) would continue to hold under various departures from actual circumstances. To the extent that we find it plausible that counterfactuals relating the occurrence of the cause to the occurrence of the effect are relevant in some nontrivial way to causal claims, it is natural to focus on considerations having to do with the sensitivity of (1.1). In other words, considering the sensitivity of 1.1 (or its analogues) allows counterfactuals relating the occurrence of the cause to the occurrence of the effect (and not just counterfactuals relating the nonoccurrence of the cause to the nonoccurrence of the effect) to do real, nontrivial work in capturing aspects of causation.

I noted above that given a causal claim of form

(1.4) c caused e ,

we can assess the sensitivity of both of the associated counterfactuals (1.1) and (1.2). I will argue, however, that the sensitivity or insensitivity of the counterfactual (1.1) typically carries more weight than the sensitivity or insensitivity of the counterfactual (1.2) in our assessment of the sensitivity of (1.4) and hence in determining our reactions to it. A causal claim of form (1.4) for which the counterfactual (1.1) is highly sensitive will, *ceteris paribus*, strike us as sensitive and hence nonstandard, problematic, or at least different from paradigmatic causal claims even if the counterfactual (1.2) is relatively insensitive. By contrast, if (1.1) is insensitive and (1.2) is sensitive, (1.4) will strike us as unproblematic. In other words, other things being equal, we value causal relationships for which

1. I make this assumption to simplify the exposition and because it is widely accepted. It is problematic, particularly in indeterministic contexts (suppose (c) is a toss of a genuinely indeterministic coin and (e) the event of its coming up heads), but nothing will turn on its correctness in what follows.

counterfactuals relating the occurrence of the cause to the occurrence of the effect are relatively insensitive, but care less about the sensitivity of the counterfactual relating the nonoccurrence of the cause to the nonoccurrence of the effect. This is not to say, however, that we do not care at all about the sensitivity of the latter counterfactual. In particular, as we shall see, if (1.1) is relatively sensitive and (1.2) is sensitive too, this may make us even more inclined to judge that the associated causal claim (1.4) is sensitive than if (1.1) is sensitive and (1.2) is insensitive. But it is the sensitivity of (1.1) that plays the primary role; the sensitivity of (1.2) is secondary.

2. Sensitivity Illustrated

To further illustrate the notion of sensitivity, let us return to the example of the vase. The sensitivity of the counterfactual

- (2.1) If the event of the rock thrown by Suzy striking the vase were to occur, the shattering of the vase would occur

has to do with whether (2.1) or its analogues would continue to hold under circumstances that differ in various ways from the actual circumstances. Put slightly differently, what we are interested in is whether (and which) counterfactuals of the form

- (2.2) If the rock thrown by Suzy were to strike the bottle in circumstances B_i different from the actual circumstances, the bottle would (still) shatter

are true for various B_i .

Some of the circumstances B_i for which claims of form (2.2) are true are so obvious that they will seem trivial. If Suzy's rock strikes the vase in Boston at the moment at which someone sneezes in Chicago, then presumably if that person had not sneezed but the world had remained relevantly similar in other respects, the bottle still would have shattered. Similarly, if we vary the color of Suzy's blouse or the price of tea in China at the time of the impact.

Other claims about the insensitivity of counterfactual (2.1) are more interesting. Readers who found plausible my claim that in the original example the timing of the shattering is counterfactually dependent on the timing of the impact (and in turn on the timing of Suzy's throw) presumably did so because they thought that if Suzy had thrown the rock (and the impact had occurred) slightly earlier or later, the counterfactual (2.1) (or its analogue) would still be true. That is, they assumed that

if Suzy had thrown the rock a bit earlier or later but in other respects the circumstances remained the same, this wouldn't affect whether (2.1) were true but only the time at which the shattering occurred. Similarly, if both the bottle and the throw had been displaced exactly ten feet to the right, then barring the presence of obstacles or other complications in the new situation, (2.1) or some analogue to it would still be true. Similarly also for some range of possible variation in temperature, wind conditions, and so on.

The judgments just described are closely related to—indeed, they are at least in part motivated or supported by—similar patterns of insensitivity that hold at the type level. Of course, when thrown rocks strike bottles, the bottles do not always shatter. Nonetheless, impacts of thrown rocks on bottles will be followed by shattering in many different circumstances—at different times, different spatial locations, and for a wide variety of other variations in background conditions.

Indeed, it is at least in part because of this that my original description seemed so natural—my claim that the impact of the rock caused the shattering seemed unsurprising in part because the introduction of such impacts will be followed by shatterings in many different circumstances. If I had instead said that the shattering was caused by Suzy's scratching her nose, this would have been more puzzling. There are certainly possible circumstances in which bottle shatterings are counterfactually dependent on nose scratchings, but these circumstances are rare and rather specialized. Typically, they involve considerable stage setting—for example, Billy promises to throw a rock at the bottle if and only if Suzy scratches her nose; Suzy scratches, Billy keeps his promise, throws, hits the bottle, and the impact causes shattering. This dependence (and, more specifically, the counterfactual relating Suzy's scratching to the shattering) would be disrupted if the stage setting were altered—for example, if Billy were not present or if no promise had been made. In contrast to the connection between the impact of Suzy's rock and the shattering, counterfactuals connecting nose scratchings to bottle shattering are likely to be rather sensitive, and this seems connected to our puzzlement when the shattering is attributed (without further elucidation) to Suzy's nose scratching.² We make the sensitive dependence between Suzy's nose

2. The anonymous referee notes that our puzzlement is also due in part to the sheer unfamiliarity of this causal connection. I agree but would note also that (i) this unfamiliarity is itself due in part to the rarity of such connections, which in turn are due in part to their sensitivity and (ii), as I explain below, my claim is that consider-

scratching and the shattering more intelligible by breaking it down, as I just did, into intermediate links of dependence that are individually less sensitive—the existence of the promise between Billy and Suzy, the link between Billy's throwing and the shattering, and so on.³

Both counterfactuals that express the dependence of effects on their causes and, on the other hand, facts about the sensitivity of these counterfactuals under changes in background conditions are relevant to understanding causation, but they correspond to different strands or elements in that notion. As one might expect, dependence counterfactuals like (1.1)–(1.2) as well as influence counterfactuals are connected to the intuition that at least in many cases, causes make a difference for their effects. Although simple versions of this intuition break down in complex cases involving preemption and overdetermination, it is nonetheless true that in many paradigmatic cases of causation, the intuition is correct—causes *are* difference makers.⁴ (Witness the rock-throwing example with which I began.) Moreover, as I and others (Halpern and Pearl 2000, Hitchcock 2001, Woodward 2003) have argued elsewhere, a more subtle version of the idea that causes are difference makers is defensible even in contexts in which there is preemption or overdetermination: causes are difference makers in such contexts when other causes or potential causes are appropriately controlled for. By way of contrast, facts about sensitivity/insensitivity are connected to the idea that causal relationships should exhibit some degree of generalizability or exportability or context independence, to ideas about the intrinsicness of causal relationships and, as we shall see, to one interpretation of the idea that causes should necessitate their effects.⁵

ations of sensitivity play *an* important role in our assessment of this and other examples, not that such considerations are the only thing that influences our assessment.

3. That is, assuming that Billy is, like most of us, the sort of person who fulfills promises over some range of conditions that is not too specialized, the link between Billy's promise and his throwing will at least not be as sensitive as the overall link between Suzy's scratching and the bottle's breaking. Similarly for the other individual links in the causal chain leading from the scratching to the breaking. Readers who are inclined to think that the sensitivity of the scratching-shattering link does not have to do with whether it is causal but rather with its role in explanation are directed to the paragraph at the end of this section.

4. For discussion of some empirical psychological evidence that people think of causes as difference makers, see Woodward forthcoming b.

5. Although I lack the space for detailed discussion, these two features of causal claims are connected in interesting and subtle ways. In particular, the strategy of dealing with cases of preemption and overdetermination by controlling for alternative

In many paradigmatic cases of causation, such as Suzy shattering the bottle, both elements—difference making/dependence and insensitivity—are present to a high degree. Nonetheless, there are also many examples in which these two elements can come apart or vary somewhat independently of each other. Recently a great deal of attention has been focused on cases, such as those involving causal overdetermination or preemption, in which counterfactual dependence fails to be present at least in its simplest or most obvious form, but in which (it is arguable) more complex forms of counterfactual dependence, exhibiting a substantial degree of insensitivity, are present when appropriate other causes are controlled for. Cases with the opposite profile—straightforward counterfactual dependence but relative sensitivity—have received considerably less attention, but they are equally interesting and important.

Most of the examples that I will explore below are of this second sort. They are examples in which there is counterfactual dependence of (what is usually taken to be) the right sort for causation (the counterfactual dependence in question is of a nonbacktracking sort, involving what I (Woodward 2003) have elsewhere called interventionist counterfactuals and so on) but in which (at least) counterfactuals of form (1.1), relating the occurrence of the putative cause to the occurrence of the putative effect (and hence the causal claims themselves), are sensitive. As we shall see, people generally respond to this sensitivity by thinking that the examples are somehow importantly different from examples involving relatively insensitive causation. This general reaction encompasses a range of more specific responses. Some may think that the relationships in at least some of the examples described below are not causal at all and hence that counterfactual dependence (which is present in the examples) is not, contrary to what is often thought, sufficient for causation. Others may think instead that although the claim that the relationships are causal is literally true for at least some of the examples, this claim is misleading or inappropriate in some way, at least if offered without additional qualification or elucidation. Still others may be tempted

causes seems to require assumptions about the sensitivity of the causal relationships involved. For example, if we hold that in the case in which c_1 and c_2 are symmetric overdetermining causes of e , c_1 counts as cause of e because, if we were to remove c_2 from the scene, e would be counterfactually dependent on c_1 —a strategy followed (with different bells and whistles in order to avoid problems we need not go into) in Woodward 2003 and Hall 2004, then we are clearly relying on an assumption of some kind about how c_1 will act in the “same way” with respect to e in both the presence and the absence of c_2 . This is an assumption about the insensitivity of the c_1 - e relationship.

to describe the examples as involving relationships that are genuinely causal but will then add that these causal relationships are shallow or unilluminating or defective from the point of view of explanation or that they lack some additional feature having to do with control or necessitation. My aim here is *not* to argue that one of these reactions is *the* correct reaction to all the examples and that all other reactions are misguided. Instead, I will content myself with urging that these reactions show that there is something distinctive about the examples and that this has to do (at least in part) with where they fall on the dimension of sensitivity/insensitivity. I should also emphasize, moreover, that as the parenthetical remark in the last sentence suggests, I do not mean to claim that sensitivity is the *only* feature influencing our reactions to the examples but merely that it plays an important (and relatively neglected) role.

3 Background and Stage Setting

Before proceeding, however, some stage setting is required. First, a remark on types and tokens: The rock-throwing example with which I began is a token causal claim, and, as I have suggested, its sensitivity has to do primarily with whether the counterfactual (2.1) would continue to hold under various sorts of changes. But one can also ask, analogously, about the sensitivity of type causal relationships or causal generalizations and the counterfactuals associated with them. For example, given the generalization

- (3.1) Alleles of type G cause phenotypic traits P ,

one can consider counterfactuals like

- (3.2) If alleles of type G were to be present in organisms of type O , then O will have phenotypic trait P

and ask whether such counterfactuals would continue to hold under various sorts of changes in O and its environment that are consistent with the occurrence of G . One can also ask analogous questions about generalizations describing relationships between variables that can take many different values (for example, that are real valued) rather than just the binary values of “present” and “absent” and about generalizations relating particular values of such variables:

- (3.3) If the pressure $P = p$ and the volume $V = v$ for a particular sample of n moles of gas, then its temperature $T = t = pv/nR$.

Here there is, of course, no unique value of P that represents the nonoccurrence of $P = p$ and similarly for the other variables. So instead of a single counterfactual along the lines of (1.2), specifying what would happen under the nonoccurrence of $P = p$, there will be a whole family of such counterfactuals, prescribing what the value of T would be for each possible alternative value of P and V . One can ask about the sensitivity of each of these counterfactuals.

Second, a remark about the kinds of causal and counterfactual claims about which we typically make sensitivity judgments: These are claims that are incomplete in the sense that they describe some but not all of the factors that are causally relevant to some outcome. For example, the claim of interest may be that C s cause E s, where there are other factors B_i in addition to C that causally influence E . Claims that are incomplete in this way are ubiquitous both in ordinary life and in the special sciences. It is of course true that if one were to replace the original claims with new claims specifying the precise relationship between C , the relevant additional factors B_i , and E , these new more precise claims (if accurate) would be highly insensitive. Very often, however, we are not able (or for other reasons do not think it worthwhile) to formulate these more precise claims and operate instead with the original incomplete claims, augmented by judgments of relative sensitivity. (Often, as in the rock-throwing example, we understand that other factors are relevant and know something about their likely distribution and influence and hence have enough information to formulate judgments of sensitivity without being in a position to formulate exceptionless generalizations.) My emphasis on the role of sensitivity in assessing causal and counterfactual claims is not meant to deny that it is desirable to formulate more precise claims when it is possible and worthwhile to do so but, rather, to describe practices that, so to speak, come into their own when we deal with incomplete generalizations. A Laplacean intelligence with knowledge of all of the fundamental laws of nature, full information regarding initial conditions throughout the universe, and no calculational limitations would have little use for the notion of sensitivity (at least as I have characterized it), but needless to say, we are not in this position.⁶

6. A related point is that the sensitivity of an incomplete generalization or an incomplete token causal claim will depend in part on how specific and fine-grained their formulation is. If the generalization specifies that if Suzy throws a rock with exactly such and such momentum, then the bottle will shatter in exactly such and such a way (or if the causal claim is that Suzy's throwing caused the bottle to shatter in exactly

Lest this last remark be misunderstood, let me add that although an important part of the motivation for focusing on causally incomplete claims has to do with our epistemic limitations, this does not mean that judgments of sensitivity are themselves knowledge-relative. Given claims of the form *Cs cause Es* (or if *C* were the case, then *E* would be the case), whether such claims would continue to hold under variations in background conditions is not itself dependent on anyone's knowledge.

Next, a remark about the relationship between sensitivity and invariance: I (Woodward 2003) have argued elsewhere that a necessary condition for a functional relationship linking (type of event or property) *C* to (type of event or property) *E* to be causal is that the relationship be invariant (in the sense that it would remain stable or continue to hold) under changes of a special sort in *C* produced by a special sort of causal process called an intervention. Sensitivity is a different although related notion: sensitivity has to do with whether causal claims or counterfactuals linking the occurrence of *C* to the occurrence *E* would continue to hold under changes in *other conditions* *B_i*, distinct from *C* and *E*. In order to have a convenient label for these conditions *B_i* I will call them background conditions.

What sorts of changes in background conditions matter for the assessment of sensitivity? At a very general level, the assessment has to do with whether the relationship of interest (understood either as a causal or a counterfactual claim) would continue to hold under changes that do not depart too much from the actual state of affairs or that do not seem too far-fetched or that are not judged to be unimportant or irrelevant for subject-matter-specific reasons. This in turn encompasses a range of more specific cases. One straightforward possibility is that we have a theory describing a state space representing a range of alternatives to the actual state of affairs and an associated metric specifying the distance between the actual state of affairs and these alternatives. Consider a driver who comes upon an icy patch in the road, begins to skid, loses control of the car, and then regains it, ending up on the shoulder of the road and not in the ditch beyond it. Among the alternatives to

such and such a way), then these claims, even if true in the actual circumstances, will be relatively sensitive, assuming that under small changes in background conditions, the same throwing would lead to a slightly different shattering. By way of contrast, more coarse-grained claims linking, say, any throwing that strikes the bottle to the bottle's shattering (in some way or other) may be less sensitive. For additional discussion of the role of coarse graining in causal judgment, see Woodward forthcoming a.

the actual state of affairs are those in which the driver hits the icy patch at a different speed or from a different angle. If his actual speed is 60 mph, an alternative in which he is traveling 55 mph is naturally viewed as closer to actuality than alternatives in which he is traveling 30 mph or 90 mph. Similarly, if the driver has had three glasses of wine, an alternative in which he has two is closer to actuality than an alternative in which he has none. If the outcome of the encounter between the driver and the ice would have been very different under small variations in the driver's speed or alcohol consumption—if the driver would have gone into the ditch if he had been traveling at 60 mph or had 2.5 glasses of wine, then the causal relationship between hitting the ice patch and the outcome is relatively sensitive.

In other cases, while there may be considerable agreement about whether some departure from actuality is substantial, this judgment may rest on considerations that are multifaceted and perhaps more difficult to make precise. If the actual situation is one in which Suzy throws a rock and it travels via an unimpeded path toward the bottle, striking and breaking it, most people would agree that alternatives in which there is a thick, solid steel barrier between the rock and the bottle or in which a second person is present who throws a rock that deflects Suzy's rock in flight represent relatively large departures from actuality, at least in comparison with alternatives in which the rock Suzy throws still strikes the bottle but with a slightly different momentum. If, were Suzy to throw, the bottle would shatter in alternative scenarios in which the momentum of the throw is different in various ways, in which environmental conditions (temperature, wind speed) vary, and so on, we are likely to think of the connection between the throw and the shattering as relatively insensitive, even though it is true that under other more far-fetched variations (the presence of the steel barrier, and so forth), this counterfactual would no longer hold.

Many considerations may influence such assessments of far-fetchedness: how improbable or uncommon the changes in question seem, either in general or given the particular circumstances of the example, whether the changes require alterations in background structures or institutions that are ordinarily stable and/or difficult to change, and so on. In addition, it also seems plausible that there is an asymmetry, at least in everyday thought, between, on the one hand, *adding* new background circumstances to the actual situation and, on the other, *removing* circumstances that are present in the actual situation—the former generally strikes us as more of a departure from actuality or at least as

less appropriate for the assessment of sensitivity.⁷ In the nose-scratching case, at least part of the reason why we regard the causation as sensitive is that the bottle shattering would not have occurred in the absence of Billy's promise. By contrast, we are considerably less likely to judge that the causal relation between Suzy's throw and the shattering is sensitive on the grounds that had the steel barrier been present, the throw would not have led to shattering. We view the former change as a matter of removing something (Billy's promise) from the original situation, rather than adding something and hence as a smaller or more allowable change. By contrast, the insertion of the barrier is viewed as adding to the original situation and hence as a larger change or at least as less appropriate for the assessment of sensitivity.

In still other cases, the specific sorts of changes that are regarded as particularly important for the assessment of sensitivity may depend on subject matter or disciplinary specific considerations. As I (Woodward 2003) have noted elsewhere, many generalizations in economics would be disrupted by surgical or pharmacological changes that alter fundamental neurological processes in economic agents—the generalizations are highly sensitive to such changes. However, this sort of sensitivity is usually not regarded as interesting or important by economists. By contrast, the failure of an economics generalization to be stable under changes in the information available to (psychologically normal) economic agents or under changes in relative costs is typically regarded as much more significant and is often taken to indicate that the generalization cannot play a fundamental or foundational role in economic explanation. In effect, for the purposes of doing economics, “abnormal” changes in neurological processing are taken to involve irrelevant or ignorable departures from actuality while changes in information or relative costs are regarded as highly relevant to the assessment of stability.

Similarly, as I shall argue below, if a generalization relating some phenotypical or behavioral trait to its supposed genetic causes is highly sensitive to various sorts of external environmental changes (especially changes in social practices that are “imposed” by others), this will often lead us to regard the claim about genetic causation as potentially misleading: degree of sensitivity under such environmental changes is particularly important in connection with generalizations about genetic causes, although it may be much less significant in connection with

7. The distinction between additions and removals is closely linked to the distinction between presences and absences that is discussed in more detail in section 5.

generalizations that have to do with other sorts of subject matters. Yet another example, also discussed below, has to do with causal claims involving the behavior of human beings. Here we tend to regard the sensitivity of such claims and the counterfactuals associated with them to changes in human action and choice as particularly important, while sensitivity under changes in other factors, not involving choice and action, carries less weight.⁸

So far I have been focusing on the extent to which (counterfactual or causal) relationships hold under this or that specific departure from actuality. However, there is also a closely related but more generic question that it sometimes makes sense to ask: given some specified set of possible background conditions associated with the relationship and some measure over this set, does the relationship hold under many or most such background conditions or only for a few? In other words, given, say, that the counterfactual

(3.4.) If C were the case, then E would be the case

is true in the actual background circumstances B_A , are counterfactuals of the form (3.4) true for most alternative background circumstances B_i for $i \neq A$ in some reference set? To the extent that the answer to this question is “yes,” the counterfactual (3.4) will count as relatively insensitive. Suppose, for example, that a certain genotype G_I leads to the phenotype P_I when an organism is at an average temperature between 15°C and 18°C during the first week of prenatal development, but that G_I leads to phenotypes different from P_I for all other temperatures during this period consistent with the survival of the organism. Particularly if average temperatures lower than 15°C or higher than 18°C are common in the organism’s environment, it will be natural to think of the relationship between G_I and P_I as relatively sensitive or at least as more sensitive than a genotype/phenotype link that holds for a much wider range of temperatures.

As I have characterized them, judgments of sensitivity depend on a number of factors that are highly context dependent. Moreover, a number of these factors may seem “subjective,” at least in the sense that they are influenced by particular interests and purposes that human beings

8. The notion of “cognitive impenetrability” as described in Pylyshyn 1984—roughly the requirement that fundamental generalizations in cognitive science be insensitive to changes in a subject’s beliefs and desires—represents yet another example of a subject-matter-specific sensitivity condition.

have. Does this undermine the usefulness of the notion of sensitivity? For a variety of reasons, I do not think so. First, as emphasized above, one of my primary interests in the role of sensitivity is in using this notion to describe actual practices of causal judgment. It is an empirical question to what extent people's judgments of sensitivity depend on the factors I have described; to the extent that they do, it is not an objection to the account that some of these features strike us as "subjective."⁹ Second, it is also of course an empirical question to what extent there is intersubjective agreement in people's judgments of sensitivity; it may be that we are largely able to agree on such judgments despite their highly contextual and highly multifaceted character. Finally, one obvious response to worries about subjectivity and context dependence is to relativize judgments of sensitivity to particular sets of changes in backgrounds. Even if you and I disagree about whether such and such a departure B^* from actuality is large or far-fetched, it may be an "objective" matter (or at least a matter about which we may expect far more agreement) whether some counterfactual or causal claim would hold under B^* . Thus, even if we disagree about whether the introduction of a solid steel barrier between Suzy and the bottle represents a large departure from the actual state of affairs, we can presumably agree that if such a barrier were introduced, it would no longer be true that if Suzy were to throw, the bottle would shatter.

Duke University Press/Journals

Next some brief remarks about the relationship between sensitivity and some other features that a generalization may possess. First, insensitivity is not the same as exceptionlessness. A generalization might hold without exception (given the actual circumstances) and yet be relatively sensitive or fragile in the sense that it would fail to hold under even very small departures from the actual circumstances. Of course, if a generalization is highly sensitive, it will often be reasonable to expect that (sooner or later) some of the situations in which it would fail to hold will be among those that are or will become actual; if so, the generalization will have exceptions. However, nothing in the notion of sensitivity requires this. Conversely, a generalization might have a few exceptions which occur only under conditions that are rare or far-fetched but may

9. More generally, it seems to me that it is a mistake to *begin* with the nonnegotiable premise or intuition that causation or causal judgment must be "objective" in this or that respect and to use this as a basis for rejecting accounts that do not respect this intuition. The respects in which causation is or is not objective ought to emerge as a conclusion of one's analysis. See Woodward 2003, chap. 2 for a more extended discussion.

hold for all or almost all background circumstances that do not represent large departures from actuality; if so, the generalization will be relatively insensitive despite having (or at least admitting the possibility of) exceptions. The second law of thermodynamics has exceptions but is nonetheless highly insensitive for typical systems containing a large number of molecules—given a mole of gas, the second law will hold for almost all initial assignments of position and momentum to the molecules in such a system.

Second, is the sensitivity of the causal relationship between *C* and *E* simply a matter of the “length” of (or the “number” of intermediate steps in) the causal chain from *C* to *E*? No. For one thing, length and number of intermediate steps (like the notion of the “directness” of a causal relationship) are clearly relative to how we grain things or to the choice of vocabulary.¹⁰ With respect to a list confined to variables like “shooting” and “death,” the causal link between shooting and death looks short and direct; with respect to an expanded list of variables that might be used to describe the chain of physiological changes in the victim that eventuate in death (damage to the heart, loss of its ability to pump blood, failure of oxygen to reach the brain), the causal chain from shooting to death looks longer and less direct, but including these intermediate variables does not make the causal relationship between shooting and death any less sensitive. As several of the examples described below will illustrate, one may have causal chains that seem (intuitively) relatively long that are fairly insensitive as well as chains that are “short” (for example, many chains involving omissions) that are sensitive.

Finally, it will have occurred to many readers that the notion of insensitivity bears at least a family resemblance to conditions imposed on causal claims by other writers. Although a detailed comparison must be beyond the scope of this article, some very brief remarks may help to provide some orientation. In the econometrics and structural-equations literature, the requirement is often imposed that equations representing causal relationships (or at least “deep” or “important” causal relationships) should be “autonomous” or “structural,” where this means that those equations should represent relationships that are stable or robust under changes that may occur elsewhere in the system being modeled, which is just to require a certain kind of insensitivity in those relationships. In the literature on probabilistic theories of causation, a so-called unanimity condition is often imposed: the cause must raise the prob-

10. For additional discussion, see *ibid.*

ability of its effect across “all” background conditions in some population of interest. Again this condition is naturally viewed as a (very specific and strong) insensitivity requirement. There is also a connection between insensitivity and the idea that the causal relation is “intrinsic.” There are many ways of spelling out the notion of “intrinsicness,” but one natural construal is that part of what is meant is that the existence of a causal relationship between c and e should not depend on (should not be sensitive to) happenings that seem in some relevant sense external to the c - e relationship or only adventitiously or accidentally related to it, with there being, of course, many different ways of spelling out what “external” means.¹¹ As we will see, the violation of this expectation in many cases of causation by omission or double prevention is connected to our judgment that the causation present in these cases is somehow nonparadigmatic.

4. Lewis on Insensitive Causation

Lewis (1986: 186) explains the notion of sensitive causation as follows:

When an effect depends counterfactually on a cause, in general it will depend on much else as well. If the cause had occurred but other circumstances had been different, the effect would not have occurred. To the extent this is so, the dependence is sensitive.

Duke University Press/Journals

Lewis’s example of sensitive causation is his writing a strong letter of recommendation that causes someone—call him X —to get a job he would not otherwise have gotten, which in turn causes someone else, Y , who would have gotten the job in the absence of Lewis’s letter, to take another job instead, and so on. The letter thus causes X , Y , and many others, who are similarly displaced from jobs they otherwise would have taken, to meet and marry different people and to have children they would not have had in the absence of the letter. Thus as a consequence of the letter some people exist and eventually die who otherwise would not have existed. So their deaths are counterfactually dependent and hence (since on Lewis’s view counterfactual dependence is sufficient for causation) caused by the recommendation. Lewis (*ibid.*) describes this as “comparatively sensitive causation” because, as he puts it, “there are many differences that would have deflected the chain of events. But if

11. For example, Lewis (1986) suggests that “whether a process going on in a specific region is causal depends only on the intrinsic character of the process itself and on the relevant laws.”

you shoot your victim point blank, only some very remarkable circumstances would prevent his death.” This second action—the shooting—involves “insensitive” causation.

Assume that Lewis writes the letter, and let N be some person who will exist in the future, eventually dying, who would not have existed if Lewis had not written the letter. Both of the counterfactuals

(4.1) If Lewis were to write the letter, N would die

and

(4.2) If Lewis were to not write the letter, N would not die

are true. Thus N 's death is counterfactually dependent on and, on Lewis's (1986) view, caused by the letter writing. Consider first the counterfactual (4.1). This counterfactual is highly sensitive: if Lewis had written the letter, but the actual circumstances had been different in various ways that are not at all far-fetched, N would not have existed and hence would not have died. For example, if contrary to actual fact, N 's great-grandfather Z had not decided to move to city A where the job was located that N 's great-grandmother X got as a consequence of Lewis's letter, or if Z had lingered a little less long in the bar in A where he met X , or if the department at University B had not decided to offer Y a position with the result that the position at A became available to X and so on, then N would not have existed. Similarly for many other changes. By contrast, (4.2) is far less sensitive. In most scenarios in which Lewis does not write the letter of recommendation, N does not exist and hence does not die. This illustrates my earlier claim that it is the sensitivity of the counterfactual linking the occurrence of the cause to the occurrence of the effect, rather than the sensitivity of the counterfactual linking the nonoccurrence of the cause to the nonoccurrence of the effect, that matters most for the overall sensitivity of the causal claim.

As Lewis says, the sensitivity present in the letter of recommendation example contrasts with the insensitivity present in his shooting example: There is a clear intuitive sense in which a shooting through the heart at close range with a large caliber bullet that causes death is not similarly sensitive to variations in background circumstances. Although one can perhaps imagine circumstances in which such a shooting would not cause death—for example, the victim is in a state-of-the-art hospital with the world's best heart surgeon who is ready to implant an artificial heart—such circumstances are very uncommon and seem far-fetched. For some very large range of ordinary variations in background circum-

stances, the shot would have continued to cause death. Again, this is reflected in the insensitivity of the counterfactual linking the occurrence of the shooting to the death.

Lewis goes on to link the difference between insensitive and sensitive causation of a death to the difference between killing someone and merely causing him to die: killing is causing to die by insensitive causation. I will return briefly to this suggestion at the end of this article, but the primary point I want to make is that the contrast Lewis draws seems to mark a real and important distinction regardless of whether one agrees with Lewis about the semantics of “kill.”

The letter of recommendation example also illustrates several other general claims made earlier. First, it shows that one can have a relationship of counterfactual dependence (involving nonbacktracking counterfactuals) even though the associated causal relationship is relatively sensitive (where again the primary test for the sensitivity has to do with whether the counterfactual relating the occurrence of the cause to the occurrence of the effect would hold under various changes). Whether or not a relationship is one of counterfactual dependence and whether or not the associated counterfactuals are sensitive or insensitive are thus distinct questions or distinct dimensions of assessment. Second, putting aside for the moment the question of what one should think, as a normative matter about the truth of the claim that

(4.3) Lewis’s writing the letter of recommendation caused the death of *N*,

it seems clear that many people will regard (4.3) as a nonstandard example of a causal relationship even if both (4.1) and (4.2) are true.¹² As remarked earlier, I think that relatively little turns, at least for the purposes of this article, on exactly what we say about the literal truth of (4.3).¹³ What I take to be clear is that the sensitivity of (4.3) leads us to regard it as importantly different from more standard cases of causation.

12. One indication: Every audience to which I have presented this material has responded to (4.3) with laughter.

13. As Chris Hitchcock has remarked to me, there is a distraction in this example. Ordinarily, when we talk of some event as causing *N*’s death, we have in mind a contrast between the actual outcome in which *N* dies and a contrasting alternative situation in which *N* exists and survives, rather than an alternative in which *N* never exists. It might be argued that this observation by itself suffices to account for whatever is odd about (4.3). However, it seems to me that (4.4), Lewis’s writing the letter of recommendation caused the existence (or the birth) of *N*, is also odd or nonparadigmatic. Those who

5. Omissions

The example just described may seem to be an isolated curiosity, but in fact once one begins to look for them, examples of relatively sensitive counterfactual dependence are quite common—indeed they include many of the examples on which recent discussion of causation has focused. In each case, the examples raise similar questions of whether “real” causation is present.

There is an ongoing debate about whether omissions or absences can be causes. On the one hand, there are many cases in which c does not occur (and no event that is relevantly like c occurs), e occurs, and in which it is true that if c had occurred, e would not have occurred—thus cases in which the occurrence of e is counterfactually dependent on the nonoccurrence of c —but in which we are reluctant to judge that the nonoccurrence of c caused e . Some theorists go further, claiming that absences or omissions are never causes, at least for what they take to be the literal or primary notion of “cause.”¹⁴ There are at least two broad motivations for this position. One appeals to the idea that there is a fundamental metaphysical distinction between presences and absences. For example, Beebe 2004 holds that causation is a relation between events and that absences, being literally “nothing” or mere privations, are not events. A second and related argument is that if c is to cause e , there must be a causal process in the sense of Salmon 1984 and Dowe 2000 connecting c to e —a spatiotemporally continuous processes that transmits energy and momentum (or perhaps some other conserved quantity) from c to e . Absences are not connected to other events via such processes and hence can be neither causes nor effects.

Although the issue deserves more detailed discussion than I can give it here, both of these considerations seem less than fully persuasive.

are attracted to the alternative diagnosis of (4.3) may substitute (4.4) instead, which is not susceptible to the alternative diagnosis.

14. See, for example, Dowe 2000. Although Dowe claims that absences cannot literally be causes, he goes on to introduce a notion that he calls quasi-causation or causation*, according to which absences can be quasi-causes of outcomes. For example, the claim that the absence of c causes the absence of e , although literally false, may be true if “causes” is glossed as “quasi-causes.” The claim that the absence of c quasi-causes the absence of e is analyzed as: if c had occurred, then c would have caused e . Among the many limitations of this analysis is the fact that it gives us no insight into why we make the discriminations we do among causal claims involving absences, that is, why we more readily regard some cases in which there is counterfactual dependence between an absence and an outcome as cases of causation than we do other such cases.

First, as an empirical matter, it is uncontroversial that people often do judge that absences and omissions are causes (and effects). More interestingly, again as an empirical matter, people discriminate among relationships of counterfactual dependence involving absences, judging that some are causal (or more paradigmatically causal) and others are not (or at least less paradigmatically causal).¹⁵ The arguments just described contend that these judgments are never correct and, whatever their normative appeal, tell us nothing about the bases of the judgments that people in fact make.

The metaphysical argument that absences cannot figure in causal relationships relies on the idea that the presence/absence distinction is a “natural” and transparent distinction that can serve as an independent input to our causal judgments, allowing us to disqualify candidate causal claims that relate absences. However, in many cases the presence/absence contrast (at least insofar as this influences the causal judgments we are prepared to accept) seems to derive from a deeper contrast between what might be described as the normal or default outcome in a specified kind of situation and departures or deviations from it, with the former being labeled “absences” and the latter “presences.” This default/deviation contrast (and the absence/presence distinction associated with it) looks contextual and theory-relative, rather than something that might plausibly be grounded in fundamental metaphysics. Indeed, it appears to be as much a *product* of our practices of causal judgment (or our practical and theoretical interests) as an independent input to them. Relatedly, many of the situations that we label as “absences” are not literally “nothing” in the way envisioned by the metaphysical account. To the extent this is so, it seems unwarranted to hold that metaphysical considerations alone rule out the possibility that absences can ever be causes.¹⁶

15. For some relevant empirical evidence, see Baron, Spranca, and Minsk 1991.

16. As an illustration of these points, consider death. Is it a presence or an absence? Looked at one way, it is about as clear a case of an absence (of life, of brain activity) as one could imagine. Nonetheless, we often treat death as a “positive” occurrence, and it is certainly not a mere “nothing” in the sense in which we might describe, say, an empty box as containing nothing. (Typically a body is present that is in a certain condition.) No one doubts that deaths can be effects (inquests investigate the causes of death), and it seems arbitrary to deny that deaths can be causes (of physiological changes to the body, of grief on the part of others, of the collapse of the empire) as well. Indeed in ordinary cases in which someone is relatively healthy and not under some immediate mortal threat, it seems odder to ask for the cause of that person’s continuing to live the following day (should he or she do so) than to ask for the cause of death, should that person die. A similar point holds for causes of death. Lack (of access to) oxygen or food

What about the claim that the existence of a causal process connecting c to e is a necessary condition for it to be true that c causes e ? It is well known that imposition of this requirement would exclude a great many causal claims that people in fact accept as true, including many claims that we ordinarily think of as involving straightforward mechanical causation, such as mechanisms that work via the release of potential energy.¹⁷ For example, the requirement has the result that releasing a billiard ball from one's hand does not cause it to fall; that if pulling the trigger of a gun releases a spring that causes (via a connecting process) the cartridge to explode, then pulling the trigger does not cause the gun to fire; that blowing up a dam does not cause the water behind it to

certainly sound like absences or privations, but we ordinarily think of these as among the causes of death, and we have “positive” ways of describing what happens when they occur (asphyxiation, starvation).

As suggested above, a natural way of thinking about such examples is in terms of a normal or default outcome and deviations from it. Given a healthy person who is not under any obvious threat, the default outcome is that he or she continues to live through the next day. Death is a deviation from this default and particularly calls for causal explanation. Deviations from the default are identified as positive events or occurrences (as changes); the realization of the default is seen as “negative”—a non-event in which nothing (relevant) changes. When there is departure from the default, whatever accounts for the contrast between the departure and the default is likely to be counted as a positive event—hence as a cause—and this may include events like asphyxiation and starvation.

We can see how the identification of a default outcome is a theory- or context-of-inquiry-relative notion, rather than a context-free metaphysical distinction, by switching the context from the commonsense one in which the default is that people go on living to the context of biological theorizing. Now we focus on the complex and intricate coordination of biochemical processes that sustain life and continued life looks like a positive, surprising event. Closely related points are made in Maudlin 2004 and Hitchcock n.d. I am indebted to their discussions.

Let me add, though, that this whole set of issues deserves far more discussion than it receives here. The anonymous referee has very plausibly insisted that some absences have distinctive features that arguably don't sit well with the claim that the presence/absence distinction is entirely context relative. For example, some absences (for example, Billy's failure to shoot down the enemy fighter in the double prevention example in section 6) seem to lack nonarbitrary spatiotemporal locations, or parts into which they are divisible. Interestingly, some of the other candidates for absences considered above do not seem to suffer from some of these limitations—for example, one can assign a time and place of occurrence to a death. Perhaps when absences lack features like a nonarbitrary spatiotemporal location and so on, this makes us think of them as relatively unsuitable candidates for causes and effects, independently of the considerations of sensitivity on which I focus.

17. For relevant discussion, see Schaffer 2000.

rush out, and so on. A massive amount of empirical evidence shows that human causal judgment is heavily influenced by what psychologists call contingency information: information about whether the effect occurs or would occur in the presence or absence of the candidate cause as well as in the presence or absence of other candidate causes.¹⁸ Information of this sort can be described in terms of counterfactual dependencies like (1.1)–(1.2) or relations among conditional probabilities. As an empirical matter, this contingency information tends to trump information about connecting processes in causal judgment. Relationships tend to be judged as causal when there is evidence for the right sort of contingency even if information about connectedness is lacking; if there is evidence that the right sort of contingency is lacking; relationships are not judged as causal even when connecting processes are present.¹⁹

I conclude from these considerations that there is no good reason to accept the claim that absences or omissions can *never* be causes. A more satisfactory position would hold that absences are sometimes causes but that there are also cases in which an outcome is counterfactually dependent on an absence but in which the outcome is not caused (or is not readily judged by us to be caused) by the absence. Or, to adopt a more cautious formulation, it is at least true that we are much more willing to accept some claims in which absences figure as causes than others. What then underlies the distinctions we make concerning the causal role of absences? Depending on the example under discussion, several different kinds of considerations may be at work, including (as a number of writers have observed) ideas having to do with moral responsibility and with whether the absences in question are viewed as “normal” or “abnor-

18. See, for example, Dickinson and Shanks 1995 and, for additional discussion, Woodward forthcoming b.

19. It is relevant in this connection to note that there is substantial empirical evidence for a dissociation between (what psychologists call) *perception* of casual relationships in collision phenomena of the sort studied by Albert Michotte and causal *judgment*. Roughly, causal judgment even regarding cases of mechanical collisions is highly sensitive to contingency information, while causal perception is not and is instead influenced by information about spatiotemporal contiguity and connecting processes. Thus the claims made above about the priority of contingency information seem true in general for casual judgment but not for causal perception. See Schottmann and Shanks 1992 for additional discussion. It is interesting to note that the empirically grounded causal perception/judgment dissociation seems to closely track the contrast in the philosophical literature between causal process theories and theories (like counterfactual accounts) that emphasize the idea that causes are difference makers.

mal.” Without wishing to deny the importance of the considerations just described, I want in what follows to explore a distinct possibility that has received little attention: that in a number of cases in which we are reluctant to say that some absence is a cause (even though the right sort of counterfactual dependence is present), an important reason for our reluctance is that the relationship between the absence and its putative effect is highly sensitive and that by contrast, the absences that we accept as causes often involve relationships that are less sensitive.

Let me begin with a contrast between two examples. First, consider

- (5.1) My writing of this very essay was caused by my not being hit by a large meteor

and the associated counterfactuals,

- (5.2) If I were not struck by a large meteor, I would have written this very essay

and

- (5.3) If I were struck by a large meteor, I would not have written this very essay.

Assume that (5.2) and (5.3) are both true. The counterfactual dependence of e_1 (= my writing this very essay) on c_1 (= my not being struck by a large meteor) makes some people tempted to say that (5.1) is literally true. On the other hand, this claim strikes others as quite counterintuitive—so much so that they regard examples like (5.1) as clear counterexamples to the contention that counterfactual dependence is sufficient for causation. Arguably, even many who think that (5.1) is literally true will also think that there are important differences between (5.1) and other more paradigmatic cases of causation—differences that they will attempt to explain in some way (for example, via appeal to pragmatic considerations) that do not undermine the truth of (5.1).

Turning to (5.2), we see that it is highly sensitive. Even if I had not been struck by a meteor, there are many relatively small changes that might have occurred under which I would not have written this essay—I might not have been invited to give various talks that prompted me to begin writing the essay, I might not have had certain conversations that have influenced the content of this essay, and so on. In most close-by neighborhoods in which I’m not hit by a meteor, I don’t write this very essay. This observation, I suggest, plays some role in people’s reluctance

to accept (5.1) or in their inclination to think that there is something about (5.1) that separates it from more paradigmatic cases of causation. In contrast, (5.3) is presumably quite insensitive. In most close-by scenarios in which I am struck by a large meteor, I die.²⁰

The asymmetry between the sensitivity of (5.2) and the insensitivity of (5.3) provides further illustration of my claim that it is the sensitivity or insensitivity of counterfactuals relating the occurrence of the putative cause (in the case of (5.1), the nonoccurrence of the meteor strike) to the occurrence of its putative effect rather than the sensitivity or non-sensitivity of the counterfactual relating the nonoccurrence of the cause to the nonoccurrence of the effect that exerts the primary influence on our judgments regarding the causal claim. (5.1) strikes us as nonparadigmatic because (5.2) is sensitive even though (5.3) is insensitive.

We can provide further support for this analysis by reversing the facts of the case. Suppose that the counterfactual claims (5.2) and (5.3) remain true but that now I *am* struck by a meteor and die, never writing this essay. Consider the claim that

(5.4) Being struck by a meteor caused me not to write this essay.

Now our focus shifts to the counterfactual (5.3) since this is now the counterfactual relating the occurrence of the putative cause in (5.4) to its effect. Since (5.3) is insensitive, we judge that there is nothing odd or nonparadigmatic about (5.4) qua causal claim, despite the fact that (5.2), relating the nonoccurrence of the cause to the nonoccurrence of the effect in (5.4), is sensitive. Again this assessment matches our intuitive judgment about (5.4).

Let us now contrast these variants on the meteor example with two other causal claims involving absences:

20. Is there some other way of explaining the problematic status of (5.1) without resorting to the claim that absences are never causes? Considerations of moral responsibility do not seem applicable. It is true that (around here, right now) strikes by large meteors are rare or abnormal, and this might prompt the thought that what underlies our reaction is some principle to the effect that nonoccurrences of events, the occurrences of which are rare or abnormal are not judged to be causes, while non-occurrences of events that are normal are at least sometimes judged to be causal. While I don't wish to deny that this consideration plays some role in our reaction to (5.1), I doubt that it is the whole story. It seems to me that (5.1) would be judged problematic by most speakers even if deadly meteor strikes were common. This suggests considerations of sensitivity play some independent role in our reaction to (5.1) even if other considerations are at work as well.

(5.5) Absence (of access) to oxygen caused *N*'s death.

(5.6) Many German civilians were caused to die from starvation (that is, from absence of food) by the British naval blockade of 1919.²¹

I take it to be uncontroversial that we are much more likely to regard (5.5) and (5.6) as acceptable than (5.1)—indeed many people will regard (5.5) and (5.6) as unproblematic, despite the fact that they seem to attribute causal efficacy to absences and despite the fact that there is no connecting causal process running from the putative cause to its effect.²² Certainly when an agent acts so as to bring about the absence described in (5.5)—I place a plastic bag over your head or remove all of the oxygen from an airtight room into which you have wandered, and death results, we seem to have little hesitancy in thinking of these as acts that cause death. Similarly for the blockade described in (5.6) or if I were to lock you in a cell and refuse to feed you. I trace the greater acceptability of (5.5)–(5.6) in comparison with (5.1) to the relative insensitivity of counterfactuals like

(5.7) If *N* were to lack access to oxygen, *N* would die,

and

(5.8) If civilians were denied access to food, they would die.

Under most non-far-fetched cases in which background conditions vary from the actual circumstances, people who are deprived of oxygen or food die.

Next let us contrast all of the cases above with yet another case, also involving an absence or omission: *Y* is a critically ill patient whose care is the exclusive responsibility of Doctor *Z*. The nature of *Y*'s disease is such that he will die if he is not given various medications and no one

21. This disgraceful incident, which occurred after the cessation of hostilities, is described in Glover 1999, 64–68. I do not deny that the judgment that the British had a moral responsibility not to interfere with supply of food to Germany influences our reaction to (5.6). On the other hand, it seems plausible that this judgment of moral responsibility is itself influenced by our assessment of the causal structure of this situation. Were the connection between the blockade and death more sensitive, we would hold the British less responsible for those deaths, as is illustrated by the famine relief example below.

22. One might of course claim that the absence of oxygen and the absence of food are positive events rather than genuine absences, but as I remark in note 15, this just raises the question of how we are to identify absences.

else is permitted to administer the medications if *Z* doesn't. *Z* fails to administer the medications and *Y* dies.

Here the relevant counterfactuals are

(5.9) If *Z* omits to administer the medication, *Y* dies,

and

(5.10) If *Z* does administer the medication, *Y* does not die,

and the causal claim of interest is:

(5.11) *Z*'s omission to administer the medication caused *Y*'s death.

How sensitive is the counterfactual (5.9)? On the one hand, it looks considerably less sensitive than the counterfactual (5.2), relating the failure of the meteor to strike me to my writing the essay. Given *Z*'s omission, *Y*'s death will occur under a significant range of non-far-fetched scenarios, as long as we continue to suppose that the disease is present in the same form and no one else besides *Z* is able to administer the medicine. Moreover, (5.10) is also (we may assume) relatively insensitive, and this may lend some additional, but secondary, support to the judgment that (5.11) is relatively insensitive. My suggestion is that these facts help to explain why we are considerably less reluctant to accept (5.11) than to accept (5.1).

While (5.11) is less sensitive than (5.1), there are nonetheless variations in background conditions that, although they depart in some significant ways from the actual situation, do not seem at all widely far-fetched or implausible under which (5.11) (and (5.9)) would no longer hold. Most obviously, it is not at all a far-fetched possibility that *Y*'s disease might have been absent or present only in a milder, nonlethal form, and if so, *Y* would not have died, even given *Z*'s omission. (This contrasts with, for example, (5.5), where it is not a serious possibility that the physiological conditions that in conjunction with the absence of oxygen are sufficient for death might be absent in a normal human being.) Similarly, *Y* would have survived even given *Z*'s omission, if someone else had administered the medication. In many medical contexts, this is not just possible but likely—usually hospitals are set up in such a way that others will become aware of *Y*'s condition and will administer the medicine if *Z* does not. In both these respects, (5.11) differs from (5.5) and (5.6). The upshot is that while (5.11) is less sensitive than (5.1), it is more sensitive than (5.5) and (5.6). This is reflected, I believe, in our reaction to these claims—(5.5) and (5.6) strike us as more paradigmatic-

cally causal (more readily acceptable, and so on) than (5.11), which is in turn less problematic than (5.1).

For purposes of additional comparison, consider a final example. Suppose that you omit to send \$10 to a famine relief organization *O*, *X* who lives in country *A* dies of starvation, and *X* would have lived if you had sent the money. Letting *c* be your omission to send the money and *e* *X*'s death, the relevant counterfactuals are:

(5.12) If *c* were to occur, then *e* would occur,

and

(5.13) If *c* had not occurred, then *e* would not have occurred.

Thus *e* counterfactually depends on the nonoccurrence of *c*. As before, I take it that while some people may hold that

(5.14) Your omission to send money caused *X*'s death,

others will be reluctant to accept this claim, at least without some further qualification. The causal claim (5.14) will strike most people as at least different from paradigmatic, true, causal claims. Again it is natural to trace these reactions, at least in part, to the fact that (5.12) is rather sensitive. Even supposing that you omit to send the money, if any one of a number of things had happened differently—if the corrupt dictator who runs country *A* had stolen a little less foreign aid from other sources, if the food transportation network in *A* had not been disrupted by war, if *X* had not been weakened by previous malnutrition, and so on—*X* would not have died.

Moreover, in this case, unlike the other examples in this section, it is also plausible that the counterfactual (5.13) linking the nonoccurrence of the cause and effect is also quite sensitive. Even if (5.13) is true (as we are assuming), it is very plausible that if background circumstances had been different in various ways, then even if you had sent the money, *X* would still have died. This would happen, for example, under small variations in the behavior of organization *O* itself (they decide to spend a bit more money in country *B* and less in country *A*) or in the food distribution network within *A*. This sensitivity of (5.13) may contribute in a secondary way to the overall judgment that (5.14) is sensitive. The sensitivity of (5.13) also has the consequence that in the variant case in which you do send the money to *O* and *X* survives, the causal link between your sending the money and *X*'s survival is sensitive: it is more

like the link between Lewis's letter of recommendation and the deaths it "causes" than the link between shooting someone point-blank and the death this causes. Again, this seems to reflect our intuitive judgment: the dual sensitivity of both of the counterfactuals (5.12) and (5.13) is associated with the judgment that the causal claims connecting the omission to *X*'s death and sending the money to *X*'s survival are both sensitive.

Before leaving the topic of the causal role of absences and omissions, a more general point is worth underscoring. It is a feature of common sense causal thinking that we distinguish between what Patricia Cheng (1997) has called *generative* and *preventive* causes. There is an asymmetry between these two kinds of causes. Roughly speaking, the operation of a preventing cause requires the presence of a generative cause but not vice-versa. Preventive causes prevent by interfering with or blocking generative causes; when no relevant generative cause is present, a candidate preventing cause, even if present, is not viewed as operative or efficacious. Thus, if someone has ingested poison (a generative cause of death), the ingestion of an antidote may prevent death, but if no poison has been ingested, then although someone who takes the antidote will (let us assume) not die, we do not think of this as a case in which the antidote acts as a preventer of death, or as a cause of survival, presumably because the outcome of survival would be the same regardless of whether the antidote is ingested. As Cheng shows, there is considerable empirical evidence that ordinary subjects make causal judgments in accord with the generative/preventive distinction.

What is true of prevention also seems true of causation by absence (or omission): we think of it as parasitic on the presence of some additional generative cause, presumably because causation by absence usually or always involves the nonoccurrence of some cause that would have been a preventer if it had occurred. For example, in the example described above, the role of the doctor's omission in causing the patient's death clearly depends on the presence of some additional generative cause (the presence of an infection or disease of some kind) that will produce death if not interfered with. As noted above, this has the consequence that the counterfactual linking the omission to the patient's death will at least be sensitive to changes that remove (or sufficiently modify) this generative cause—in the absence of the lethal form of the disease (we assume) the patient's death will not occur even if the doctor omits to administer the medication. More generally, this fact—that the causal efficacy, such as it is, of omissions depends on the presence of an

additional generative cause—means that there is an important respect in which causation by omission will *always* be somewhat sensitive: omissions will no longer cause their effects under variations in which the relevant generative process is no longer present (or sufficiently modified), and this will show up in the fact that the counterfactual relating the omission to its effect is sensitive to such changes.

At this point, the reader may well wonder about how this analysis applies to (5.5) and (5.6). What are the generative causes present in these examples? In fact, such causes *are* present—they are just hard to see. In both cases, the generative causes are those physiological processes P that in the presence of food and oxygen sustain normal life and in their absence produce toxic effects—cell death, and so on. But while the presence of the lethal disease strikes us as an adventitious and readily modifiable feature of the situation in (5.11), the processes P are “normal” features, the removal or relevant modification of which would be very far-fetched. We have little idea, for example, of what sorts of changes in P would be required for a human being to lack access to oxygen for an extended period of time and yet survive and no conception of how to bring about such changes—at the very least they would have to be “extraordinary.”

Finally, let me draw attention to a very interesting connection between sensitivity and Lewis’s notion of influence. Suppose that the causal relationship between X and Y is relatively insensitive in the sense that in the presence of X , Y would continue to occur at a time, place, and manner determined by X , under changes in other causes of (including background conditions that are causally relevant to) Y . Then these other causes or background conditions will have relatively little influence in Lewis’s sense on Y , at least in the presence of X . Put the other way around, suppose that X and Z are both causes of Y and that Z has a lot of influence on Y in the sense that changes in the value of Z , its time and place of occurrence, and so forth are associated with changes in the value, timing, and so on of Y . Then X is likely to be a relatively sensitive cause of Y . This is reflected in many of the examples discussed above. Shooting someone through the heart at point-blank range is a relatively insensitive cause of death, and this means that various other factors have relatively little influence on the time and manner of death, given that such a shooting has occurred. Similarly, suppose that my decisions concerning the content of a certain paragraph of this essay at a time t shortly before its composition at $t + d$ have a great deal of influence on

the content of that paragraph in the sense that different decisions will lead to different contents. Then my failure to be struck by a large meteor between t and $t + d$ will be, at best, a relatively sensitive cause of the content of that paragraph.

6. Double Prevention

Cases of double prevention are cases in which the occurrence of some event or process, c_1 , would prevent some outcome, e (which would occur in the absence of c_1), but in which some second event, c_2 , prevents the occurrence of the potential preventer, c_1 , thereby allowing e to occur. In such cases, the occurrence of e depends counterfactually on the occurrence of c_2 (since c_2 and e occur, and if c_2 had not occurred, c_1 would have occurred and would have prevented e). Nonetheless, many people will feel at least some resistance to the claim that c_2 causes e .

An example of Ned Hall's (2004), from whom the phrase "double prevention" is taken, illustrates the basic idea. Suzy's plane will bomb a target (e) if not prevented from doing so. An enemy pilot will shoot down Suzy's plane (c_1) unless prevented from doing so. Billy, piloting another plane, shoots down the enemy pilot (c_2), and Suzy bombs the target. Suzy's bombing counterfactually depends on whether Billy shoots down the enemy pilot. Nonetheless, some—perhaps many—find problematic the unqualified claim that

- (6.1) Billy's shooting down of the enemy pilot caused Suzy's bombing.

In this connection, Hall notes that the relationship between Billy's action and the bombing lacks several important features that (he thinks) we ordinarily associate with causation. One of these is what Hall calls "locality": Billy's shooting may occur at a great spatiotemporal distance from the bombing—hundreds of miles away and hours earlier, with no intervening events connecting the firing to the bombing in a spatiotemporally continuous way. In addition, the relationship of counterfactual dependence between Billy's shooting and the bombing itself depends on features that seem "extrinsic" to that relation. For example, in a variant on the original example in which the enemy's superiors were about to order him not to attack Suzy when Billy shot him down, Suzy's bombing would no longer depend on Billy's action, even though the variant seems "intrinsically" just like the original example, differing only in the far-off intentions of the enemy's superiors. Even if we don't share Hall's assessment, it seems hard to deny that there is *some* important difference

between the connection of Billy's action to the destruction of the target and, say, the connection of the act of shooting, in Lewis's example, to the death of the victim or, for that matter, the connection between Suzy's pressing the button that releases the bombs and the destruction of her target, even though this, like Billy's action, involves the removal of a preventer (whatever holds the bombs in place) and no transfer of energy/momentum from cause to effect.

Hall's own diagnosis is that double prevention examples show that we operate with two distinct concepts of causation. One concept (which Hall calls "dependence") involves counterfactual dependence but is not transitive and does not require a spatiotemporally continuous process connecting cause and effect or the determination of causal structure by intrinsic features. The other concept ("production") satisfies suitably formulated conditions of locality, intrinsicness, and transitivity but need not involve counterfactual dependence. Billy's firing is a cause of the bombing in the dependence sense but not in the production sense. By way of contrast, shooting someone at point-blank range or breaking a bottle with a thrown rock are paradigmatic cases of production.

It should be clear from my discussion so far that where Hall sees two distinct "concepts" of causation, I see instead a single, more unitary concept with two different but interrelated strands or elements, both of which should be understood in terms of the (different) counterfactual commitments they carry. As explained above, one of these strands has to do with the counterfactual dependence of the effect on the cause; the other, with the sensitivity of the relationship between cause and effect. What we have in many cases of double prevention, including the connection between Billy's firing and the bombing, is counterfactual dependence of a relatively sensitive sort. It is this, at least in part, that leads us to judge—to the extent that we do—that such cases lack some feature that characterizes paradigmatic cases of causation, despite the presence of counterfactual dependence. Or to put the matter another way, it is the combination of counterfactual dependence and relative sensitivity that leads many of us to have the ambivalent reactions that we do to Hall's example—the presence of counterfactual dependence leads us to judge that the case involves genuine causation; the relative sensitivity of this counterfactual dependence inclines us to deny this judgment. On my view, Hall's two distinct concepts instead represent variations along two different dimensions of a single concept, with paradigm cases of production involving the presence of a relatively insensitive causal rela-

tionship and dependence having its usual meaning of counterfactual dependence.²³

Why do I say that the connection between Billy's firing and Suzy's bombing is relatively sensitive? First, there are many possible changes—changes that do not seem at all unlikely or far-fetched—under which the counterfactual

- (6.2) If Billy were to shoot down the enemy pilot, Suzy would drop the bombs

would no longer be true. This would happen if, for example, Suzy had changed her mind about carrying out the bombing or if a second fighter who would shoot down Suzy unless interfered with in some way had been present. More generally, as is true of all cases involving prevention, the causal efficacy of Billy's preventive activities are parasitic on the presence of another, potentially generative cause (the action of the enemy fighter in shooting down Suzy), and the role of this cause, qua preventer of Suzy's bombing, were it to be effective, would be parasitic on the potentially generative cause represented by Suzy's activities. Moreover, in this case, we also have sensitivity of the counterfactual relating the nonoccurrence of the cause to the nonoccurrence of the effect:

- (6.3) If Billy were not to shoot down the enemy fighter, Suzy would not drop the bombs.

and this may reinforce our judgment about the sensitivity of (6.1). For example, if another fighter had been present who would have shot down the enemy if Billy had not, or if the enemy pilot had eluded Billy and then changed his mind about pursuing Suzy, and so on for many other possibilities, Suzy's bombing would presumably still have occurred and hence (6.3) would be false.

In addition to the sheer number of contingencies that would disrupt (6.2) and (6.3), their character matters too. As suggested above, we are inclined to attach a particular significance to the sensitivity of causal and counterfactual claims under changes in human decisions and actions. Other things being equal, to the extent that a relationship would be disrupted if human actors were to act or choose differently in

23. Space precludes detailed discussion of Hall's very interesting characterization of production. For what it is worth, however, I am inclined to think that transitivity is not a reasonable condition to impose on any concept of causation, for reasons I describe in Woodward 2003, 57–59.

any one of a large number of different ways that seem not at all unusual or far-fetched, we will be particularly inclined to regard that relationship as sensitive. In the Billy/Suzy bombing example (as well as in the aid-to-charity example), this sort of sensitivity is present in spades. If Suzy had changed her mind about the desirability or morality of the bombing, the counterfactual (6.2) would not hold. If the enemy had not decided (or had not been ordered) to fly his plane that morning, or if he had eluded Billy but then decided not to attack Suzy, the counterfactual (6.3) would be false. In a discussion of this example, Joseph Halpern and Judea Pearl (2000) contend that part of our reluctance to regard Billy's action as straightforwardly a cause of the bombing stems from our sense that this is to treat Suzy (and the enemy pilot) as automatons rather than as agents for whom there is a genuine possibility of choosing differently. They suggest that if we substitute mechanical preprogrammed drones for the enemy and Suzy in the example (so that the enemy-drone will automatically and inevitably shoot down Suzy if not intercepted by Billy and the Suzy-drone will automatically carry out the bombing unless shot down by the enemy-drone), our resistance to treating Billy as a cause of the bombing becomes less pronounced. This is an empirical claim that strikes me as probably correct. A sensitivity-based account that assigns a special importance to human actions makes good sense of this feature of causal judgment.²⁴

Let us now consider by way of contrast another system involving double prevention: the synthesis of the enzymes that metabolize lactose in *Escherichia coli*. This occurs when and only when lactose is present in the bacteria's environment. Simplifying greatly, the mechanism involved in this synthesis works in the following way: when lactose is absent, a repressor protein is synthesized that prevents the genes that would otherwise synthesize the enzymes from being transcribed; when lactose is present, the repressor is inactivated and the relevant genes contribute to the production of the enzymes. The presence of the enzymes counterfactually depends on the presence of lactose in the environment, but this is because the presence of lactose prevents the presence of something that if present would prevent the synthesis of the enzymes. In this case, needless to say, human agency plays no role in the operation of the system,

24. Hart and Honoré 1985 is a classic discussion of the sensitivity of causal attribution in the law to considerations having to do with human agency. See also Honoré 2005, especially the discussion of "later intervention" in cases involving human agents.

and it is well buffered against a range of environmental contingencies, at least of the sort that are likely to occur or involve modest departures from the actual environment. In addition, the combination of the role of lactose in the synthesis of the repressor protein and the role of the latter in turning off and on the genes involved in the synthesis of the lactose enzymes is not an accidental or ad hoc feature of this system but is, rather, a normal, “designed” feature that is due to natural selection: the presence of just one component of the system (for example, the synthesis of the repressor protein in the absence of lactose) without the other would confer no selective advantage at all.²⁵ For all of these reasons, the counterfactual dependence between lactose and the synthesis of the enzymes strikes us as relatively stable or insensitive, despite the fact that it involves double prevention and despite the fact that there is no connecting process from cause to effect.

Consistently with this, we seem to have little hesitation in concluding that this is a straightforward case of causation—indeed genetics texts describe this as a case in which the presence of lactose “controls” (but via “negative” rather than “positive” control) the presence of the enzyme. I take examples like this to suggest that it is not the presence of double prevention *per se* that makes us reluctant to regard a relationship of counterfactual dependence as causal; instead, it is the relative sensitivity of the associated counterfactuals that prompts this reaction. In addition, this example reinforces the conclusion that we reached in section 6: that relatively insensitive causation may be present even when there is no spatiotemporal connecting process or transfer of energy and momentum from cause to effect. It may well be true, as an empirical matter, that such processes are often realizers of insensitive causation, but they are not the only such realizers. Moreover, such processes may be present even when a causal relationship is highly sensitive.

25. Chris Hitchcock has observed to me that it is probably true in general that we are more willing to regard relationships of double prevention as causal when there has been natural selection for such relationships or when the relationship is present in an artifact and is the result of human design. Thus, to take an example of Jonathan Schaffer’s (2000), we have no problem with the suggestion that pulling the trigger of a gun causes it to fire even if the gun works by double prevention. Similarly, it seems unproblematic to claim that Suzy’s pushing the bomb bay button in her plane causes the bombs to fall, even if this also involves the removal of a preventer: the doors to the bomb compartment. In both cases, the overall relationship of counterfactual dependence is rather stable and designed to be so. Changes that would disrupt the dependence would need to be rather “abnormal.”

7. Genetic Causation

In his book *Rethinking Social Policy* (1992), Christopher Jencks asks us to suppose that that external sexual characteristics (having male rather than female genitals, breasts, and so forth) are genetically caused. (If the notion of genetic causation of gross phenotypical features is at all coherent, this claim is surely true.) Jencks also notes that in our society and many others, social expectations and customs based on the presence of such external sexual characteristics lead women to have, on average, different hair length than men. If one wishes to represent the relationships in this example, the natural way to do so seems to be as follows:

- (7.1) Male genotype (G_1) \rightarrow male sexual characteristics (S_1) \rightarrow Social expectations for short hair (E_1) \rightarrow short hair length (L_1).
- (7.2) Female genotype (G_2) \rightarrow female sexual characteristics (S_2) \rightarrow Social expectations for long hair (E_2) \rightarrow long hair length (L_2).

Alternatively, summarizing both (7.1) and (7.2) in terms of a single representation employing two-valued variables we have

- (7.3) Genotype (G) \rightarrow sexual characteristics (S) \rightarrow Social expectations for hair length (E) \rightarrow hair length (L),

where an arrow from X to Y means that Y is counterfactually dependent on X .²⁶

Should we think of this as establishing that

- (7.4) Having a male (female) genotype *causes* short (long) hair length

or that

- (7.5) Genotype causes hair length?

As with many of the other examples discussed above, (7.4) and (7.5) strike many of us as, if not literally false, at least misleading or problematic in some way. This is certainly the reaction that Jencks, who uses the example to raise questions about what it means to say that a trait is heritable, expects us to have to it.

26. Note that in this case, it may be that neither of the alternative forms of G (G_1 and G_2) are naturally regarded as a “presence” or a “positive” (as opposed to a negative) state. Assuming that G_1 and G_2 are the only alternatives, the counterfactual linking the nonoccurrence of G_1 to the nonoccurrence of L_1 will be just (7.2), and (7.1) will play a similar role in connection with (7.2).

Given my discussion so far, it will come as no surprise that I trace the problematic character of (7.4) and (7.5) to the fact that both are relatively sensitive or unstable and more particularly to the fact that they are unstable under changes that are regarded as particularly important in the assessment of claims of genetic causation—changes in social conventions or expectations. Obviously, if societal expectations were to change so that, for example, women were expected to wear short hair and never long hair or both sexes were expected to shave their heads, the counterfactual dependence between gender and hair length expressed by (7.4)–(7.5) would be broken. More generally, even if a phenotypic trait is counterfactually dependent on genotype, the claim that the trait is genetically caused tends to strike us as misleading or problematic if that counterfactual dependence is itself highly unstable under environmental manipulations or changes in social practices and conventions that take place well outside people’s skins. On this view of the matter (7.4) and (7.5) are problematic or potentially misleading for broadly the same reason as Lewis’s claim to have caused many deaths by writing a letter of recommendation is misleading; in both cases, we have counterfactual dependence without (much) insensitivity or at least without the kind of insensitivity that seems most relevant and important in the case at hand.

Consider another example, also due originally to Jencks (in his 1980) but much discussed in the subsequent literature (for example, Block 1995, Sesardic 2003). Assume that (natural) hair color is genetically caused and suppose that redheaded children are systematically beaten and mistreated because of their hair color. This has the result that, on average, they have lower IQs (*IQ*) than other children who are not mistreated. We thus have the following chain of counterfactual dependence:

$$(7.6) \quad \text{Genotype } (G) \rightarrow \text{Haircolor } (H) \rightarrow \text{Mistreatment } (M) \rightarrow \text{IQ}.$$

As in the previous example, the dependence of *IQ* on genotype involves nonbacktracking, interventionist counterfactuals; the dependence is not, for example, the result of the operation of some common cause of both *G* and *IQ*, not the result of a relationship in which *IQ* causes *G*, and so on. Thus, in the envisioned scenario, according to Lewis’s theory,

$$(7.7) \quad G \text{ causes } IQ.$$

Now contrast this example with another. Scientists have recently discovered a gene (the *APSM* gene) that strongly influences cortex size (Cho 2004). People with a mutated form of this gene have abnormally small cortices (microcephaly) and are seriously mentally retarded. Here

too we have a case in which there is counterfactual dependence of *IQ* on the presence or absence of a certain genetic structure (that we may represent by means of the variable *G'*)

$$(7.8) \quad G' \rightarrow IQ.$$

Virtually everyone will agree that the claim that *G'* causes *IQ*—or at least that abnormal forms of the gene cause mental retardation—is more natural and less misleading if the facts are as described in (7.8) than is the corresponding claim about the genetic causation of *IQ* if the facts are as described in (7.6). The analysis that I favor traces this difference to the different sensitivities of (7.7) and (7.8). (7.7) would be disrupted by all sorts of not-very-far-fetched changes in the social environment of red-haired children—most obviously by changes in how they are treated. By contrast, it seems very unlikely that any changes in social environment or indeed any changes that take place outside the skin of those having the abnormal form of *APSM* will disrupt (7.8). (7.7) is, in an obvious sense, much more mutable, changeable, or modifiable than (7.8), and this difference affects our reaction to it.

Jencks uses this example to make some familiar points about the notion of heritability (as measured by the so-called coefficient of heritability h^2) as used by geneticists. Coefficients of heritability have a number of features that make them misleading measures of anything like the degree to which a trait is under genetic control. Some of these features have to do with the fact that heritability coefficients, like correlation coefficients more generally, are population-relative. In different populations in which the same causal relationships hold between genotype, environment, and phenotypic trait the heritability of the trait will vary depending on how the variance of the phenotypic trait or variance of the genotype varies in those populations.²⁷ The observations that I made above do not turn on considerations of population relativity but rather have to do, as I have said, with considerations of sensitivity. Even if we focus on the causal relationships between genotype and phenotype that hold at the level of individuals rather than populations, so that heritability measures do not apply, those relationships may vary greatly in their degree of sensitivity. The problem with heritability coefficients is not just that they are population-relative; they also fail to reflect considerations of sensitivity.

27. The heritability h^2 of a trait is just the ratio of the variance of the trait that is due to genetic causes to the total variance: $h^2 = s_g^2/s_t^2$.

Interestingly, this view of matters fits naturally with a further distinction that is commonly drawn in discussions of heritability. The examples described so far in this section all involve what geneticists call *reactive* genotype-environment (*G-E*) covariance. Intuitively, this is *G-E* correlation that arises because a certain type of environment is (as it is commonly described) *imposed* on subjects with a certain genotype and this in turn influences whether they have some phenotypic trait, as when other people abuse red-haired children, lowering their *IQs*. This is distinguished from *active* *G-E* covariance that arises because a given type of environment is selected or chosen by the subjects themselves because of their genotype and this in turn influences their phenotype. Block (1995) offers the example of a child with strong musical abilities with a genetic basis who seeks out and pays attention to musically rich environments (listens to music, takes music classes, and so on), with the result that the child becomes even more musically talented. Additional examples are furnished by animals, such as beavers, who seek out and construct distinctive environments where this behavior has a genetic basis.

As a general rule, researchers are much more willing to incorporate correlations that arise because of active *G-E* covariance into estimates of heritability than they are willing to do this with reactive *G-E* covariance. That is, if there is a correlation between parental and offspring *IQ* that arises only because red-haired parents tend to have red-haired children and both are abused and have low *IQs* as a result (non-red-haired children being treated normally), this correlation is not taken into account in calculating the heritability of *IQ*. By way of contrast, it is thought to be appropriate to include any correlation between the musical abilities of parents and offspring that arises because of the active covariance associated with seeking out musically rich environments in estimates of the heritability of musical ability. Similarly, if beavers with superior dam-building abilities (that we assume have a genetic basis) are more likely as a result to live in food-rich environments and hence to be larger and more fit, their larger size is taken to be a heritable trait. The rationale for this is that when the environment is actively selected in this way, “the environmental influences can be plausibly regarded as just a way a genotype is expressed, and hence as ‘a more or less inevitable result of genotype’” (Sesardic 2003 1005, quoting Jinks and Fulker 1970: 323).

This differential treatment of reactive and active *G-E* correlation makes a good deal of sense if we see it as grounded in judgments of differential sensitivity. In general, when the counterfactual dependence of a phenotypic trait on some genotype is entirely the result of reactive *G-E*

covariance, that dependence will be relatively sensitive, as we see in the example involving the *IQs* of red-haired children. In contrast, when this counterfactual dependence arises as a result of active covariance, it will be less sensitive to normal perturbations in background circumstances. Thus in the example described above, the supposition is that in a relatively large range of normal environments, the musically talented children will successfully seek out musically rich environments and this will in turn affect ability. To the extent this is so, these environments begin to look like a “more or less inevitable result” of their genotypes rather than something that is adventitiously imposed on them. In other words, the differential treatment of active and reactive *G-E* covariance follows the general practice of thinking of relatively insensitive genotype-phenotype relations of counterfactual dependence as more paradigmatically causal, more naturally interpretable as “genetic” causation, and so on.

8. Interaction and Causal Importance

It is frequently observed that most or all phenotypic traits are the result of an “interaction” between an organism’s genotype and its environment. The moral drawn by many writers (for example, Oyama 1985) is that it makes no sense to claim that either of these causal factors is in any way more “important” or “privileged” in the production of the phenotypic trait—instead the role of both is equal and symmetrical. In fact, however, this interactionist observation is perfectly consistent with an asymmetry of sensitivity in the relationships between *G* and *P* and between *E* and *P*. Consider the genotype/environment/phenotype relationships represented in figure 1.

For each genotype and environment type, there is an interaction in the sense that each genotype produces at least two different phenotypes in different environments and each environment produces at least two different phenotypes in the presence of different genotypes. Nonetheless, if we treat each type of genotype and environment as equally “important” for the purpose of assessing sensitivity, asymmetries in sensitivity are also

	E_1	E_2	E_3	E_4
G_1	P_1	P_1	P_1	P_2
G_2	P_1	P_4	P_5	P_6
G_3	P_1	P_3	P_3	P_5
G_4	P_3	P_4	P_2	P_1

Figure 1. Gene-Environment Interaction

present. For example, G_1 produces P_1 in “most” of the environments (E_1, E_2, E_3) in which it occurs. In this sense, the relationship between G_1 and P_1 is relatively insensitive to changes in environment. By contrast, G_2 produces P_1 only in the environment E_1 and not in any other environment. In this sense, the relationship between G_2 and P_1 is relatively sensitive. A similar asymmetry is present on the environmental side, with environment E_1 producing the same phenotype P_1 outcome for most genotypes (G_1, G_2, G_3), while E_3 produces P_1 only for the genotype G_1 .

Do these asymmetries in sensitivity have any significance? Let me make two tentative suggestions. The first is that (echoing some of my remarks in section 7) *part* (not the only part but an important part) of what people who use notions like “innate” or “genetically specified” sometimes mean is that there is a relatively insensitive relationship running from some feature of the genotype to some phenotypical trait of interest—the relationship is stable across many different, even if not all, environments. Thus, for example, the phenotype P_1 is (relatively) innately specified by the genotype G_1 . If G_1 – G_4 are alternative forms (alleles) of some gene G and, instead of the situation envisioned in the table above, G_1 produced the same phenotype P_1 in almost all environments, G_2 a different phenotype P_2 in almost all environments, and so on, then we might speak of the phenotypes P_1, P_2 , and so forth as (relatively) innate and as genetically specified by G . By contrast, if, as is the case with G_3 and G_4 in the table above, the same allele produces different phenotypical traits in all or most different environments, we will be more inclined to think of that trait as “acquired” rather than innate.²⁸ For example, the suggestion that hair length is innate or genetically specified in Jencks’s example strikes us as misguided, at least in part because the relationship of counterfactual dependence between the presence of the chromosome Y and hair length is so unstable or sensitive. By contrast, eye color and external sex characteristics strike us as more plausible candidates for innately specified traits because the relationship between genotype and these traits is relatively insensitive. It is not that these relationships cannot be disrupted at all, but rather that

28. Of course if a genotype is to play a causal role at all in the production of a phenotype, there must be counterfactual dependence between some changes in the phenotype and changes in the genotype for some environment type. This condition is satisfied in the example under discussion because changes in genotype (for example, from G_1 to G_2) are associated with changes in phenotype for some (indeed most) environments. Obviously, satisfaction of this condition can coexist with either sensitivity or insensitivity of the $G_1 \rightarrow P_1$ relationship to changes in environment.

the kinds of changes that would disrupt them would need to be relatively radical and “non-environmental”—for example, changes elsewhere in the genome.²⁹

The differences in degree of sensitivity to which I have drawn attention also seem to be connected with another kind of judgment we make—judgments of causal importance. There are no doubt many different things that might legitimately be meant by the claim that one cause of an outcome is “more important” than some other cause of the same outcome. But one thing that such claims may mean is associated with asymmetries in sensitivity. If, say, causal factor A_1 will produce outcome E_1 as long as some other factor B is in any one of a large number of states $B_1 \dots, B_n$, but, for each of these states B_i , the outcome that results is highly sensitive to whether A_1, \dots, A_n is present (that is, B_1 produces one outcome when A_1 is present, a different outcome when A_2 is present, and so on and similarly for B_2), then, other things being equal, A_1 will be regarded as a “more important cause” of E than any of $B_1 \dots, B_n$.³⁰ If we were to modify the table in figure 1 so that while G_1 continued to produce P_1 in each of the environments E_1-E_3 , E_1 produced different phenotypes for each of the genotypes G_1-G_4 , and similarly for the other environments, then in this sense, G_1 would be a more important cause of P_1 than, say, E_1 .

As another illustration of the same idea, consider the debate between those who emphasize the role of natural selection (adaptation-

29. To think of “innateness” along the lines I have been describing is to think of it in terms of notions like “flatness or invariance of norms of reaction,” as Philip Kitcher (2001) recommends, or as having to do with the degree of “developmental canalization” of a trait, as Roger Ariew (1996) proposes. This is a feature that virtually all commentators agree is relevant to at least some of the judgments of innateness that people seem prepared to make—it is one of the “meanings” of innateness. Unfortunately, as is often noted (for example, Griffiths 2002), researchers associate many other features with “innateness” as well—features that neither imply nor are implied by the notion of relative insensitivity to environmental changes. My remarks above are not intended as an argument that the characterization in terms of insensitivity yields the only acceptable characterization of innateness. Nor do I want to deny that equivocation among different meanings of “innateness” has made the term highly liable to misuse. My intention is only to draw attention to one aspect of the meaning of terms like “innate” and “genetic” and to exhibit their connection with ideas about sensitivity—ideas which, as I have sought to illustrate, figure in a number of other ways in our practices of causal judgment.

30. Again, I emphasize that I am *not* claiming that the notion under discussion is the *only* legitimate notion of causal importance. No doubt much confusion will result if one conflates this notion with other possible notions.

ists) and those who emphasize the importance of developmental constraints in the evolution of various character traits. This is often framed as a debate about the relative “importance” of these two factors since for many traits it seems clear that an adequate explanation will need to assign some role to both factors. However, this way of viewing the debate is rarely accompanied by any very precise account of what “importance” means.

Here is one possible account.³¹ Someone who thinks that natural selection plays an important (the most important, and so on) role in the explanation of some trait *T* will think that there are environmental conditions *E* (posing some adaptive problem) such that natural selection will produce a trait like *T* in a large range of cases in which *E* is present—that is, over a wide range of different genotypes, possible phylogenetic histories, variations produced by noise or drift, and so on. In other words, it is assumed that the relationship between *E* and *T* is relatively insensitive across such changes. On this view, if environmental conditions are such that possessing eyes or wings would confer an adaptive advantage, then natural selection will operate in such a way that those traits will evolve for many different sorts of animals that may be genetically rather dissimilar, as long as they face those conditions. Of course there will be limits—for some animals, some traits that would be advantageous to possess will never evolve because those animals cannot produce the appropriate variations for natural selection to operate on—but these limits will be relatively broad and will not be constraining. If we were to “replay the tape” of the course of events that led some species of animal to develop wings, we would find that as long as the right selective forces were operative, wings would evolve in many cases, even as other factors having to do with genetics were varied. (The idea would be that this captures the intuitive idea that natural selection is a “powerful” force in shaping this trait and that other constraints are comparatively “weak.”)

By contrast, someone who believes that “developmental constraints” played an important role in a species’ possession of *T* and natural selection a relatively unimportant role (or at least a substantially less important role) will have a different set of beliefs about which relationships are the insensitive ones. Here the expectation is that if we were to take this species and other species that are genetically related and

31. For some broadly similar remarks, see Sterelny 1996. However, I would not endorse Sterelny’s distinction between what he calls “actual sequence” and “robust process” explanations.

vary the environmental conditions in which they lived in such a way that alternatives to T would be adaptively more advantageous in the new conditions, one would see relatively little variation in T . In other words, the relationship between T and certain features of its genotype (its *bauplan* or whatever) would be relatively insensitive to changes in selective regime and other sorts of sorts of variations. Contrary to what the adaptationist supposes, the power of natural selection to change this relationship would be relatively limited.

As before, I do not mean to suggest that this is the only possible gloss to be put on the disagreements between adaptationists and their critics, and I readily acknowledge that I've ignored many complexities and subtleties in those disagreements. On the other hand, I do claim that it is helpful and illuminating to see this disagreement as, at least in part, an empirical disagreement about the relative sensitivity of various relationships.

9. Conclusion

I conclude with some more general remarks. First, although I have argued that a concern with sensitivity informs many of our causal judgments, it is natural to wonder whether there is anything more general that might be said about why sensitivity matters to us. One consideration that is worth exploring has to do with the connection between insensitivity and (one notion of) *control*.

Consider an agent who wishes to bring about some outcome E by manipulating the state of a cause C of that outcome. Typically, whether E occurs will depend not just on the state of C but on many other circumstances besides. Think of these circumstances as capable of assuming any one of a large number of values B_1, \dots, B_n . For some of these values B_i , E will follow if C is introduced, for other values, E will not follow. Even if C is under the control of the agent, it will often not be under the agent's control which of the values B_i is realized. Furthermore, the agent may not know the value of B_i that obtains in circumstances in which she acts or whether that value is such that it will enable the occurrence of C if E is introduced or will instead frustrate it. In this sort of situation, it often will be a good strategy for the agent to look for a cause C of E that is such that the relationship between C and E is relatively insensitive to different values of B_i (or at least insensitive to values of B_i that are "close" to the actual circumstances). In the ideal case, the agent will be able to find a cause C that makes E unavoidable or inevitable—that is, a cause

C that is followed by E for all possible values B_i . Then it will not matter at all which value of B_i happens to be realized. Even if there is no such cause at hand, the agent can still look for a cause that makes E *relatively* inevitable or unavoidable—that is, a cause C that is followed by E for a large range of different values of B_i , hence a causal relationship between C and E that is relatively insensitive.

Consider the choice between a robust poison that is capable of causing death in a wide variety of circumstances—regardless of diet, medical condition, and so on—and a fragile poison that will cause death only in very special circumstances. Imagine a would-be murderer who does not know whether these special circumstances obtain and cannot control whether they obtain. The murderer might be well advised to employ the robust poison, even if, were the special circumstances present, the probability of death would be higher with the fragile poison. As this example illustrates, employing a C that is a relatively insensitive cause of E can give one a kind of *control* over whether E obtains and a kind of insurance against the possibility that nature or other agents might thwart one's plans that a more sensitive cause does not provide.

We see this feature—or its relative absence—in many of the examples discussed above. In Lewis's example, shooting someone through the heart is a good way of making his death nearly inevitable—inevitable not just in the sense that it makes death highly likely given the actual circumstances but also in the sense that death is probable for most nonfar-fetched variations on those circumstances. Shooting through the heart gives one a high degree of control over whether the death occurs—given the shooting, there is not much that nature or any other agent can do to prevent death. In this sense, shooting is a nearly foolproof way of ensuring death. By way of contrast, withholding money from a charitable organization is not a way of making someone's death nearly inevitable in this second, proof-against-changing-circumstances sense, even if, as it happens, in the actual circumstances the death is counterfactually dependent on the withholding. There is just too much else that easily might have happened and in happening would have disrupted the counterfactual link between the omission and the death. In this sense, the strategy of choosing withholding does not give one control over whether any particular death occurs. I suggest that this feature is closely bound up with Lewis's inclination to describe the shooting but not the withholding as a "killing"—"killing" seems like the right description when a causing of death in the counterfactual-dependence sense along with a high degree of control (that is, relative insensitivity) is present. For parallel reasons,

we judge that the presence of lactose in the environment of *E. Coli* “controls” whether it produces enzymes that synthesize lactose, despite the fact that the mechanism of control involves double prevention, while, by way of contrast, Billy’s activities do not control whether Suzy drops her bombs. A similar elucidation suggests itself for claims that some trait is under “genetic control.” To the extent that we value causal relationships that provide a basis for a relatively high degree of control, we should also value relatively insensitive causal relationships.

The notion of insensitivity (and the related notion of invariance) can also be used to provide a natural and nonmetaphysically mysterious account of what might be meant by the claim that causes “necessitate” their effects: to say that *C* necessitates *E* full stop is simply to say that *E* always occurs when *C* occurs, regardless of what else happens—in other words, that *E* is (in this sense) inevitable or unavoidable given *C* and that the relationship between *C* and *E* is insensitive or immutable under all (physically possible) changes. One might think that fundamental natural laws have this feature or something close to it. Moreover, this notion of necessitation can be relativized: the relationship between *C* and *E* is relatively necessary to the extent that there are relatively few nonfarfetched changes that might disrupt it—that is, to the extent that it is relatively insensitive. Paradigmatic causal relationships are expected to exhibit some substantial degree of immutability or necessitation: relationships like those between genotype and hair length or between Lewis’s letter and the particular deaths with which it is associated strike us as nonstandard cases of causation precisely because they lack this feature of immutability or necessitation to a very substantial extent.

Next, a remark about pragmatics. Some readers will be inclined to think that the examples I have described should be understood not in terms of the notion of variations in sensitivity/insensitivity but rather in terms of considerations having to do with the *pragmatics* of causal judgment—in particular, in terms of the fact that we find certain causal relationships more interesting or salient than others. Thus, it may be suggested that the reason why we think there is something odd or nonstandard about the claim that Lewis’s letter of recommendation caused many deaths or the claim that genotype causes hair length is that although these claims are literally true, they describe causal relationships that for various pragmatic reasons we regard as less important or which we are less likely to notice or pay attention to than others.

I am happy to endorse part of this suggestion. I agree that we do find the causal (or, if you prefer, counterfactual dependence) relation-

ships just described less important or salient than others. My remarks about insensitivity can be viewed as an explanation of *why* this is the case—we find the relationships uninteresting and nonsalient *because* they are relatively sensitive. What I want to resist is the stronger claim that the appeal to pragmatics shows that the ideas about sensitivity I have been describing are mistaken or unnecessary in assessing causal claims.

The idea that the appeal to pragmatics undermines rather than complements an account that appeals to considerations of sensitivity requires, I believe, something like the following picture: There is really no relevant objective or structural difference between, on the one hand, the relationship connecting Lewis's letter of recommendation to the deaths that follow and, on the other, the relationship that would hold between his shooting someone point-blank and that person's death. As far as what is really "out there" in the world, so to speak, both relationships are equal. Given that this is the case, the apparent difference between the relationships can only be a reflection of psychological facts about us—in particular, our tendency to find some such relationships more salient or interesting than others. Just as (it might be supposed) the amusingness of a joke is not an intrinsic feature of the joke but rather has to do with our reaction to it, so also for our different attitudes toward the examples discussed in this essay.

I have agreed that in some cases "subjective" considerations guided by human interests help to influence our judgments of how far-fetched or important various changes in background conditions are taken to be for the purposes of assessing sensitivity, although I also argued that more objective considerations matter too. But once a given change or set of changes are specified, there is every reason to think that it is an objective matter whether various counterfactuals will continue to hold under such changes. In this sense, the sensitivity of a causal claim under this or that particular set of changes is a fact about the world and not something that is made true or false just by our attitudes or interests. My claim is that it is such facts about sensitivity and not just an undifferentiated disposition to find certain relationships more interesting than others that lead us to have the differential reactions we do to the examples considered in this article.

It thus seems to me that if the role I have assigned to sensitivity is to be *replaced* with a more general appeal to pragmatics, some independent argument is required that would show either that judgments about sensitivity are not grounded (at least in part) in facts about the world or that they are not related to causal judgment in the way that I have

claimed. The invocation of pragmatics does not by itself provide such an argument since, as I have just argued, the insensitivity-based account can be viewed as a story about the basis for the judgments about importance and salience to which the enthusiasts for pragmatics draw attention. In other words, the idea that the examples of more and less sensitive causation described in this paper differ *only* in our attitudes to them and in no other way is a view we should adopt only after we have first explored more objectivist accounts of the differences between those examples and have found them wanting. This paper may be viewed as a preliminary exploration of such an objectivist account.

Finally, a remark about the goal of philosophical accounts of causation. For understandable reasons, most such accounts have tended to look for conditions that are, as it were, minimally sufficient for a relationship to qualify as causal. That is, the focus has been largely on finding necessary and sufficient conditions that distinguish causal from non-causal (for example, merely correlational) relationships. Given this orientation, issues about the status of, say, omissions, inevitably present themselves as questions about whether omissions are causes at all. On which side of the cause/noncause border do they lie?

It is worth explicitly emphasizing that there is no particular reason why philosophical theories of causation should focus *only* on questions of this sort. As I have sought to illustrate, there are also important distinctions to be drawn among claims that qualify (or that we may wish to regard as qualifying) as minimally causal. A philosophical theory of causation doesn't have to *stop* once it has distinguished the causal from the noncausal. Part of the discomfort we feel when confronted with the question of whether various sorts of omissions can be causes is that locating them on either side of the cause/noncause border can seem arbitrary and potentially misleading. On the one hand, it is undeniable that many relationships of counterfactual dependence involving omissions seem very different from paradigmatic causal relationships; on the other, if, as often seems to be the case, the contrast that is intended when we label a relationship as causal is with relationships that are "merely correlational," then there is an obvious motivation for regarding at least some omissions as causal. Rather than focusing so exclusively on the question of whether relationships involving omissions, double prevention, and so on are minimally causal, perhaps a more (or at least equally) illuminating strategy is to ask how they resemble and differ from other, more paradigmatic cases of causation. Sensitivity has, I have argued, a central role to play in answering this question.

References

- Ariew, A. 1996. "Innateness and Canalization." *Philosophy of Science* 63 (supp.), pt. 1: S19–S27.
- Baron, J., M. Spranca, and E. Minsk. 1991. "Omission and Commission in Judgment and Choice." *Journal of Experimental Social Psychology* 27: 76–105.
- Beebe, H. 2004. "Causing and Nothingness." In *Causation and Counterfactuals*, ed. J. Collins, N. Hall, and L. Paul, 291–308. Cambridge, MA: MIT Press.
- Block, N. 1995. "How Heritability Misleads about Race." *Cognition* 56: 99–128.
- Cheng, P. W. 1997. "From Covariation to Causation: A Causal Power Theory." *Psychological Review* 104: 367–405.
- Cho, A., ed. 2004. "Brain Gain Explained?" *Science* 303: 36.
- Dickinson, A., and Shanks, D. 1995. "Instrumental Action and Causal Representation." In *Causal Cognition: A Multidisciplinary Debate*, ed. D. Sperber, D. Premack, and A. Premack. Oxford: Oxford University Press.
- Dowe, P. 2000. *Physical Causation*. Cambridge: Cambridge University Press.
- Glover, J. 1999. *Humanity: A Moral History of the Twentieth Century*. New Haven, CT: Yale University Press.
- Griffiths, P. 2002. "What is Innateness?" *Monist* 85: 70–85.
- Hall, N. 2004. "Two Concepts of Causation." In *Causation and Counterfactuals*, ed. J. Collins, N. Hall, and L. Paul, 225–76. Cambridge, MA: MIT Press.
- Halpern, J., and J. Pearl. 2000. *Causes and Explanations: A Structural Model Approach; Technical Report R-266*. Los Angeles: Cognitive Systems Laboratory, University of California.
- Hart, H. L. A., and T. Honoré. 1985. *Causation in the Law*. 2nd ed. Oxford: Clarendon.
- Hitchcock, C. 2001. "The Intrasensitivity of Causation Revealed in Equations and Graphs." *Journal of Philosophy* 98: 273–99.
- Honoré, A. M. 2005. "Causation in the Law." In *Stanford Encyclopedia of Philosophy*, plato.stanford.edu/entries/causation-law (accessed November 28, 2005).
- Jencks, C. 1980. "Hereditry, Environment, and Public Policy Reconsidered." *American Sociological Review* 45: 723–36.
- . 1992. *Rethinking Social Policy*. Cambridge, MA: Harvard University Press.
- Jinks, J. L., and D. W. Fulker. 1970. "Comparison of the Biometrical Genetical, MAVA, and Classical Approaches to the Analysis of Human Behavior." *Psychological Bulletin* 73: 311–49.
- Kitcher, P. 2001. "Battling the Undead: How (and How Not) to Resist Genetic Determinism." In *Thinking about Evolution: Historical, Philosophical, and Political Perspectives*, ed. R. Singh, C. Krimbas, D. Paul, and J. Beatty, 396–414. Cambridge: Cambridge University Press.

- Lewis, D. 1986. "Postscript C to 'Causation': (Insensitive Causation)." In *Philosophical Papers*, vol. 2, 184–88. Oxford: Oxford University Press.
- . 2000. "Causation as Influence," *Journal of Philosophy* 97: 182–97.
- Maudlin, T. 2004. "Causation, Counterfactuals, and the Third Factor." In *Causation and Counterfactuals*, ed. J. Collins, N. Hall, and L. Paul, 419–43. Cambridge, MA: MIT Press.
- Oyama, S. 1985. *The Ontogeny of Information*. Cambridge, MA: MIT Press.
- Pylyshyn, Z. 1984. *Computation and Cognition: Toward a Foundation for Cognitive Science*. Cambridge, MA: MIT Press.
- Salmon, W. 1984. *Scientific Explanation and the Causal Structure of the World*. Princeton, NJ: Princeton University Press.
- Schaffer, J. 2000. "Causation by Disconnection." *Philosophy of Science* 67: 285–300.
- Schottmann, A., and D. Shanks. 1992. "Evidence for a Distinction between Judged and Perceived Causality." *Quarterly Journal of Experimental Psychology* 44A: 321–42.
- Sesardic, N. 2003. "Heritability and Indirect Causation." *Philosophy of Science* 70 (PSA 02, vol. 1): 1002–14.
- Sterelny, K. 1996. "Explanatory Pluralism in Evolutionary Biology." *Biology and Philosophy* 11: 193–214.
- Woodward, J. 2003. *Making Things Happen: A Theory of Causal Explanation*. New York: Oxford University Press.
- . Forthcoming a. "Causation with a Human Face." In *Causation and the Constitution of Reality: Russell's Republic Revisited*, ed. H. Price and R. Corry. Oxford: Oxford University Press.
- . Forthcoming b. "Interventionist Theories of Causation in Psychological Perspective." In *Causal Learning: Psychology, Philosophy, and Computation*, ed. A. Gopnik and L. Schulz. New York: Oxford University Press.