Midwest Studies in Philosophy, XXXI (2007)

The Folk Probably Don't Think What You Think They Think: Experiments on Causation by Absence

JONATHAN LIVENGOOD AND EDOUARD MACHERY

1. INTRODUCTION

Folk theories—untutored people's (often implicit) theories about various features of the world—have been fashionable objects of inquiry in psychology for almost two decades now (e.g., Hirschfeld and Gelman 1994), and more recently they have been of interest in experimental philosophy (Nichols 2004). Folk theories of psychology, physics, biology, and ethics have all come under investigation. Folk metaphysics, however, has not been as extensively studied. That so little is known about folk metaphysics is unfortunate for (at least) two reasons. First, folk metaphysics is almost certainly implicit, and it is likely to be our default way of thinking about metaphysical problems. Moreover, one's metaphysical commitments can have profound consequences—in scientific, religious, and ethical contexts, for example. Thus, folk metaphysics ought to be dragged out into the open and exposed to criticism. As Peirce eloquently remarked (1994, 1.129; see also 1994, 7.579),

> Every man of us has a metaphysics, and has to have one; and it will influence his life greatly. Far better, then, that that metaphysics should be criticized and not be allowed to run loose.

Second, folk metaphysics ought to be studied, because metaphysicians often assume without evidence that they know what the folk think, and these assumptions are sometimes wrong in important ways. In the absence of experimental evidence about folk intuitions, metaphysical speculation is *nominally* constrained

by so-called ordinary notions of identity, free will, and the like, while *really* it is checked only by the often peculiar intuitions of metaphysicians themselves.

So, are the folk monists or dualists? How do they think about identity and individuation? Are they compatibilist or incompatibilist about free will? Armchair speculation is not adequate to answering these questions, as we learned again in the course of the studies presented in this article. Fortunately, psychologists and experimental philosophers have made a start toward understanding certain parts of folk metaphysics.[1] In the present article, we consider another part, not much discussed by experimental philosophers—the folk theory of causation.[2] Although the psychological literature on causal cognition is already quite extensive,[3] we feel that psychologists have neglected some significant aspects of the folk theory of causation.

We take our cue from a recent philosophical debate about the metaphysics of causation. Philosophers disagree about whether causation is a relation between events. Some philosophers, notably Hugh Mellor (1995, 2004) and David Lewis (2000, 2004), contend that causation is *not* a relation between events; indeed, they claim that causation is not a relation in the strict sense at all. The main argument for this view is that since the folk routinely talk of absences as causes and effects, while absences, being nothing at all, cannot be the relata of any relation, denying that causation is a relation is the best way to preserve folk intuitions about causation. Not all philosophers have been swayed by this consideration, however. Helen Beebee (2004) proposes to save the view that causes are relations between events by explaining that the ordinary practice of talking about absences as causes is not due to nonrelationist metaphysical commitments, but to the subtle confusion of causal explanation with causation. According to Beebee, people assert causal sentences—an event *c causing* the occurrence of an event *e*—in place of explanatory sentences—the occurrence of *e being explained* by reference to *c*. When people say that an absence *a* causes an event *e*, their assertion is literally false, since absences—being nothing at all—can't be causes. However, such an assertion also conveys a (possibly true) judgment that the occurrence of the event *e* is explained by the absence *a*. In this respect, people's mistaken judgments are innocuous.

We thought Beebee's approach promising enough to be put to the test. We experimentally tested two claims. First, we tested the claim that the folk deny the causal status of some absences that count as genuine causes according to nonrelationist theories of causation. Second, we tested the claim that the folk really fail to distinguish causation from causal explanation. The first prediction was upheld, but, surprisingly, the second was not, suggesting that *pace* Beebee, the folk do not conflate causation and causal explanation. We argue that our results stand as a challenge to *both* Lewis *and* Beebee.

1. See, for example, Bloom (2004) on children's dualism; Rips et al. (2006) on the folk conception of temporal identity; and Nahmias et al. (2006) on intuitions about free will.

2. But see Knobe (2005) and Knobe and Fraser (forthcoming).

3. See, for example, Cheng (1997), Goldvarg and Johnson-Laird (2001), Gopnik et al. (2004), Sloman (2005), Mandel (2005), Spellman et al. (2005), and Wolff (2007).

In section 2, we describe the metaphysicians' debate as it stands at present. In section 3, we present and discuss our predictions. In section 4, we describe three experiments we conducted, present the results, and discuss their implications. In section 5, we argue that our findings challenge Beebee's account of folk judgments of causation by absence. In section 6, we argue that our findings also challenge Lewis's interpretation of these judgments. In section 7, to conclude, we examine the implications of our findings for the debate between relationists and nonrelationists as well as some future lines of research suggested by the experiments discussed in this article.

2. CAUSATION AND ABSENCE

2.1 Relationist versus Nonrelationist Views of Causation

One way of dividing up metaphysical views about the nature of causation turns on the question of whether causation is a relation between events. Call those metaphysicians who say that causation *is* a relation between events *relationists*. Following Beebee (2004), we call those metaphysicians who say that causation is *not* essentially a relation between events *nonrelationists*.[4] Davidson (1967) is an example of the first sort of metaphysician; Mellor (1995, 2004) and Lewis (2000, 2004) are examples of the second.

Suppose that a green billiard ball rolls toward a motionless black billiard ball. When the green ball hits the black ball, the latter starts to roll. Here, the impact *caused* the motion of the black ball. For a relationist theory, this is to say that the event of the impact and the event of the black ball moving (or beginning to move) are connected by the causal relation, whatever that is. The relationist view does not itself provide an account of the causal relation, which might be given in terms of counterfactual dependence (Lewis 1973, reprinted in Lewis 1986a), or transfer of a mark (Salmon 1984), or in some other way. For the nonrelationist, on the other hand, *no relation* between events captures what is really essential to causation.

Both sides usually agree that the relationist view of causation is appealing at first glance. This may be owing to a sense (correct or not) that ordinary causal discourse is best described by the relationist view. But whatever the reason, if the relationist view is initially appealing, why would anyone give it up (as Lewis in fact did)? One reason to give up the relationist view of causation is the puzzle of causation by absence.

2.2 The Puzzle of Causation by Absence

Ordinary discourse about causation often involves absences. In at least some of these cases, absences seem to appear as causes, effects, or both. When the putative

4. Some nonrelationists grant that causation is sometimes a relation between events, while insisting that being a relation between events is not essential to causation (for further discussion, see Beebee 2004 and Lewis 2004).

cause is an absence, we have a case of *causation by absence*. Whenever the putative effect is an absence, we have a case of *prevention*.

Consider a case of causation by absence. Suppose a pharmaceutical company sells a drug treating acute pain due to arthritis, but has neglected to test whether long-term consumption of the drug increases people's risk of heart attack. Suppose also that this drug in fact increases people's risk of heart attack and that several people die of heart attack after its long-term use. One might then say,

(i) The pharmaceutical company's negligence caused some people to die.

In this example, an absence—the company's negligence—causes an event—the death of some people.

Consider now a case of prevention. Suppose that Chris Duffy, the Pittsburgh Pirates' center fielder, makes a spectacular catch that saves a run. One might say,

(ii) The catch prevented a run from scoring.

In this example, an event—the center fielder's catch—causes an absence—no run scores.

Relationists deny that (i) and (ii) refer to instances of causation. Since absences are nothing at all, they cannot be the relata of any relation. Particularly, they cannot be the relata of a causal relation. Thus, if commonsense judgments about prevention and about causation by absence—such as (i) and (ii)—are literally true, relationist views of causation are incorrect. This consideration moved Lewis (2000, 2004) to give up on relationist views of causation.

However, as Lewis himself points out, nonrelationist views of causation are saddled with a symmetric problem: If there is any causation by absence, then there is an awful lot of it. But people typically don't mention the various and sundry events the absence of which preserves the status quo. No one mentions the absence of bullets, of mustard gas, of explosions, of massive amounts of radiation, of divine retribution (smiting), and so forth, all of which could make one's life both nasty and short. Moreover, one need not turn murderous in order to uncover alleged causes that no one mentions. Suppose, for instance, that during his last world tour, Mick Jagger's plants in his Londonian apartment die because they were not watered. (Only unintentional plantslaughter at worst!) Now, if G. W. Bush had watered Mick Jagger's plants, they would not have died. Thus, on nonrelationist views, G. W. Bush's failure to water Mick Jagger's plants—an absence—caused them to die. Suppose now that we ask people to explain the death of Jagger's plants. Odds are that people would *not* spontaneously utter (iii):

(iii) G. W. Bush's not watering Mick Jagger's plants caused them to die.

Instead, they would look closer to home. Perhaps it is the fault of Jagger's house cleaner or even of Jagger himself (say, because he forgot to tell his house cleaner to water the plants). If one supposes that people's failure to mention these sorts of causation by absence is tantamount to judging that such alleged causes are not

truly causes and if such judgments are correct, then nonrelationist views of causation are false.[5]

Thus, folk discourse about prevention and causation by absence raises a puzzle both for the proponents of nonrelationist views of causation and for the proponents of relationist views of causation. Nonrelationists need to explain why people fail to mention causes by absence like (iii). Relationists need to explain away folk judgments of prevention and of causation by absence, such as (i) and (ii). This puzzle is well recognized by metaphysicians. Thus, Lewis writes (2000, 196),

> The foe of causation by absences owes us an explanation of why we sometimes do say that an absence caused something. The friend of causation by absences owes us an explanation of why we sometimes refuse to say that an absence caused something, even when we have just the right pattern of dependence.

2.3 Lewis's and Beebee's Solutions

Now, the puzzle of causation by absence would more accurately be called the puzzle of *folk judgments* of causation by absence, though it is too much to hope that this clunky designation will gain widespread acceptance. After all, the problem is not so much to explain what causation by absence could be as to see whether ordinary talk about causation, which sometimes appears to treat absences as causes, commits us to believing that absences really can be causes.

Of course, ordinary talk about causation matters only when it truly expresses people's beliefs about causation. The problem is that these beliefs cannot always be read off directly from ordinary discourse about causation. Rather, such discourse may be the result of the interplay between people's beliefs about causation and conversational pragmatics. Thus, when a metaphysical theory conflicts with folk discourse, the theory may be saved by providing a plausible account of the pragmatics of the relevant discourse that explains discrepancies between this theory and folk discourse. Let a conflict between a metaphysical theory and folk discourse be called a *puzzle*, and call a solution to a puzzle *pragmatic* if it reconciles them by providing a pragmatic account of ordinary discourse. Both Lewis and Beebee provide pragmatic solutions to the puzzle of causation by absence.[6]

Lewis (2000) argues that the failure to mention numerous causes by absence *is not* tantamount to judging that these alleged causes by absence are not truly causes. Although judging that not-p is one reason for not saying that p, it is not the only one. We also do not say that p when p is not relevant, despite judging that p. As Lewis puts it (2000, 196),

5. Lewis only goes so far as to say that the folk often refrain from making true assertions about causation by absence. This is certainly a problem for the nonrelationist, but things get worse, for not only is it likely that people would not spontaneously utter (iii)—except perhaps as a joke—on hearing that Mick Jagger's plants died from lack of water, it is likely that people would *deny* that Bush's failure was a cause at all. We return to this point below.

6. Here "pragmatics" is understood broadly to cover not only an account of assertion along Grice's (1975) (or one of his followers') line, but also to include any account of assertion that refers to speakers' communicative intentions.

> [T]here are ever so many reasons why it might be inappropriate to say something true. It might be irrelevant to the conversation, it might convey a false hint, it might be known already to all concerned. . . .

For instance, when we (JL and EM) step inside the department of History and Philosophy of Science at the University of Pittsburgh, we judge, at least implicitly, that the walls are beige. But, we never say, "The walls are beige." Since everyone can see that the walls of the department are beige, there is no reason to mention this fact. Similarly, according to Lewis, people do not mention the vast majority of causes by absence countenanced by nonrelationists—for example, G. W. Bush's failure to water Mick Jagger's plants—because they are not as relevant as other causes by absence—for example, the failure of Mick Jagger's house cleaner to water his plants. As Lewis notes, his solution to the puzzle of causation by absence is not available to the relationist, because the folk do assert that some absences are causes. Thus, he claims that the nonrelationist is better able to solve the puzzle.

Beebee (2004) disputes Lewis's solution, claiming that people do not merely fail to assert that some absences are causes. Rather, she maintains that people explicitly *deny* that some absences are causes. She recognizes that this is an empirical claim, but while pointing out that she has no reason to doubt it, she does not provide any experimental evidence to support it either. She writes (2004, 301, our emphasis),

> The plants would have survived, for example, if Flora's neighbor had installed a sprinkler system that then got activated accidentally while she was away, or if her roof had started leaking at a point just above the orchids during a rainstorm (. . .), or . . . I do not think that most people would happily accept that the failure of each of these events was equally a cause of the orchids' death. *Of course, this is an empirical claim about what people are ordinarily inclined to judge.* But I have not come across any evidence to suggest that the claim is false.

If Beebee is correct to say that people deny that some absences thought to be causes by nonrelationists really are causes, then both relationists and nonrelationists are forced to reject some folk intuitions as metaphysically incorrect. Relationists must claim that people are just wrong to say that absences ever cause anything, while nonrelationists must claim that people are just wrong to say that some absences are not causes. This leaves both sides in about equally bad shape. Thus, she writes, "The moral is that commonsense intuitions about causation are no more damaging to the relationist than they are to the non-relationist" (2004, 294).

Having thus recast the puzzle of causation by absence, Beebee offers her own pragmatic solution. She explains both assertions that some absences are causes and denials that other absences are causes by hypothesizing that people make two mistakes: First, people confuse causal explanation with causation; and second, they do not distinguish between (1) sentences that are true but inadequate as explanations and (2) sentences that are straightforwardly false. Let us consider these mistakes one at a time.

Beebee draws a distinction between causation on the one hand and causal explanation on the other. Following Davidson (1967), she contends that causation requires that two events stand in the appropriate relation to one another. However, she follows Lewis (1986b) in thinking that causal explanation requires only that some information be given about the causal history of an event. Thus, a sentence asserting that an absence *a* caused an event *e* would be false. However, a corresponding sentence asserting that an event *e* occurred because an event *c* failed to occur—that is, a sentence asserting that *e* obtained because an absence *a* obtained—might be a true causal explanation. Such a sentence is true whenever it is counterfactually true that *c* would have brought about some event *e′* distinct from *e*. Hence, according to Beebee, both

(iii) G. W. Bush's not watering Mick Jagger's plants caused them to die.

and

(iv) Keith Richards's not watering Mick Jagger's plants caused them to die.

are false, but the corresponding sentences

(iii*) Mick Jagger's plants died because G. W. Bush did not water them.

and

(iv*) Mick Jagger's plants died because Keith Richards did not water them.

are both true causal explanations, since it is true that had Jagger's plants been watered—regardless of who watered them—they would not have died.

But while (iii*) and (iv*) are both true, they may not be equally good explanations. And that leads us to Beebee's distinction between (1) sentences that are true but inadequate as explanations and (2) sentences that are straightforwardly false. She writes (2004, 307):

> As a denier of causation by absence, what do I say about the commonsense distinction between "Flora's failure to water the orchids caused their death" (true) and "your failure to water the orchids caused their death" (false)? I say, of course, that both are false. But the corresponding *explanatory* claims ("the orchids died because Flora failed to water them"; "the orchids died because *you* failed to water them") are both true. However, they do not, in most contexts, count as equally *adequate* explanations. When we explain why the orchids died, we must, if our explanation is to count as adequate as well as true, be sensitive to why the explanation was requested in the first place. In the context where my interlocutor is requesting information that is relevant to the issue of whom to *blame* for the death of the orchids, it would be highly misleading of me to say that they died because you didn't water them.

> The *truth* of my utterance doesn't depend on the moral question of who is to blame; but the *adequacy* of my explanation does.

One can easily imagine a context in which (iv*) is an adequate explanation and (iii*) is not. Suppose, for instance, that Mick Jagger asked his fellow band member Keith Richards to look after his plants while he was on vacation. Unluckily for Jagger, Richards was in a drunken stupor when he agreed to do the favor, and when he finally sobered up, he didn't remember the conversation at all. Moreover, G. W. Bush was, not surprisingly, never contacted by Jagger in regard to his plants. Now suppose that on returning from his vacation, Jagger asks, "Why are my plants dead?" Under these circumstances, (iv*) is an adequate explanation, for it supplies the sort of information that is relevant to Jagger's question—that is, it tells him who is morally responsible for the death of his plants. Whereas, (iii*) is not an adequate explanation.

Beebee claims that the folk fail to draw either of her distinctions. She is explicit about both points. Thus, early in her essay she writes, "I claim that common sense judges some absences to be causes because it fails to distinguish between causation and causal explanation . . ." (2004, 293). And in the final section of her article, she claims (2004, 307),

> [T]he distinction between a true explanation—a true 'because' statement—and an adequate explanation is one that common sense has a tendency to ignore. If you judge that the orchids died because Flora didn't water them but not because you didn't water them yourself . . . you mistake lack of explanatory salience for falsity.

The upshot is that people will *assert* that an absence caused an event *e* whenever the absence provides an adequate causal explanation of *e*, and people will *deny* that an absence causes an event whenever the absence does not provide an adequate causal explanation of *e*. Assertion of causation by absence happens because people confuse causal explanation with causation, and which absence is asserted to be the cause of an event depends on what makes for an adequate explanation of that event. Denial of causation by absence happens because people confuse inadequately explanatory sentences with false ones and because people deny sentences they believe to be false.[7]

3. EXPERIMENTAL STUDY OF FOLK CAUSATION: PREDICTIONS

Our experimental study was inspired by Beebee's proposed solution to the problem of causation by absence. We derived two main predictions from her proposal. First, Beebee asserts that the folk *deny* that some absences that qualify as genuine causes according to nonrelationist theories really are causes. If this is correct, people ought to disagree with some sentences asserting that an absence caused an event. Second, Beebee explicitly contends that people conflate causation

7. For a similar reading of Beebee's proposal, see McGrath (2005, 130–31).

and explanation. If this is correct, people should agree or disagree to the same extent when given a sentence asserting that an absence *a* caused an event *e* as when given a sentence explaining the occurrence of *e* by reference to *a*.

Following some philosophers, such as Davidson (1967) and Beebee (2004), as well as most psychologists (e.g., Gopnik 1998), we assumed that sentences whose form is "*p* because *q*" express explanations. For instance, we assumed that the sentence "the black ball rolled because the white ball hit it" explains the movement of the black ball by reference to the action of the white ball on the black ball. Thus, we derived the following predictions from Beebee's account:

Prediction 1. In some contexts, people should *disagree* with sentences asserting that some absences that qualify as genuine causes according to nonrelationist theories cause events. Thus, if a seven-point scale is used to measure people's agreement with a target sentence, 1 indicating a complete disagreement and 7 indicating a complete agreement, people's mean agreement with a sentence that asserts that such an absence *a* caused an event *e* should be significantly below 4.

Prediction 2. People should agree equally with sentences asserting that an event *e* took place because an absence *a* obtained and with sentences asserting that the absence *a* caused the event *e*, when *a* is a relevant absence in the causal history of *e*.

Our experimental results were mixed. The first prediction was substantially upheld. People do deny some sentences that assert that an absence caused an event. However, the second (and arguably the most important) prediction was *not* upheld. We found that people sometimes agree significantly more with sentences asserting that an event *e* occurred because an absence *a* obtained than with sentences asserting that *a* caused the occurrence of *e*, suggesting that people actually *succeed* in drawing Beebee's distinction between causal explanation and causation. Additional findings are discussed in the next section.

4. EXPERIMENTS

4.1 The Rope Case

In order to test Predictions 1 and 2, we developed the following case:

> *Experiment 1: The Rope Case*
> "Susan had to climb a rope in gym class. Susan was a very good climber, and she climbed all the way to the rafters."

In the *causation condition*, this short description was followed by the following question:

> "On a scale of 1 to 7, 1 indicating that you totally disagree and 7 indicating that you totally agree, how much do you agree with the following claim? 'The rope not breaking caused Susan to reach the rafters.' "

In the *explanation condition*, the short description was followed by the following question:

> "On a scale of 1 to 7, 1 indicating that you totally disagree and 7 indicating that you totally agree, how much do you agree with the following claim? 'Susan reached the rafters because the rope did not break.'"

Ninety-five individuals taking classes at the University of Pittsburgh took part in the experiment (mean age: 21; 2 years; range: 18–60 years; 41.1% male). In classroom settings, subjects were randomly assigned to the causation condition (N = 49) or to the explanation condition (N = 46). Subjects had to answer the relevant question by circling a numeral on a seven-point scale, anchored at 1 with "totally disagree" and at 7 with "totally agree." Subjects were also asked to fill in a short demographic questionnaire. Three subjects answered that they were not native speakers of English and one subject failed to answer this question. Because excluding these subjects did not affect the data analysis, they were not removed from the data set.

In the causation condition, the mean answer was 2.73 (SD = 1.89), while in the explanation condition, the mean answer was 3.50 (SD = 1.71). The mean answer in the causation condition was significantly different from a neutral answer, 4 ($t(48) = -5.20$, $p < 0.001$). Subjects were also significantly more likely to agree with the target sentence in the explanation condition ("Susan reached the rafters because the rope did not break") than with the target sentence in the causation condition ("The rope not breaking caused Susan to reach the rafters") ($t(93) = 2.07$, $p < 0.05$, two-tailed; Figure 1). The effect size $d = 0.43$ corresponds approximately to a medium effect size (Cohen 1992).

Subjects disagreed with the sentence asserting that an absence (viz. the rope not breaking) caused an event (Susan reaching the rafters). Furthermore, subjects were significantly less likely to agree with a sentence asserting that an absence (viz. the rope not breaking) causes an event (Susan reaching the rafters) than with a sentence asserting that this event takes place because an absence obtains—that is, with a sentence explaining an event by means of an absence.

4.2 Control: The Broken Rope Case

To provide further support for the finding suggesting that people are more likely to agree with a sentence asserting that an event e occurred because another event e' failed to occur than with a sentence asserting that e' caused e, we controlled for two alternative explanations.

First, subjects may have found the gerundive expression "the rope not breaking" strange, because, although perfectly grammatical, it is somewhat infrequent in colloquial English. As a result, they may have agreed with the target sentence in the causation condition less than with the target sentence of the explanation condition. Second, people might be in general more likely to agree with a sentence formulated with the connective "because" than with a sentence formulated with the verb "to cause." If that were the case, subjects given the rope case would be less likely

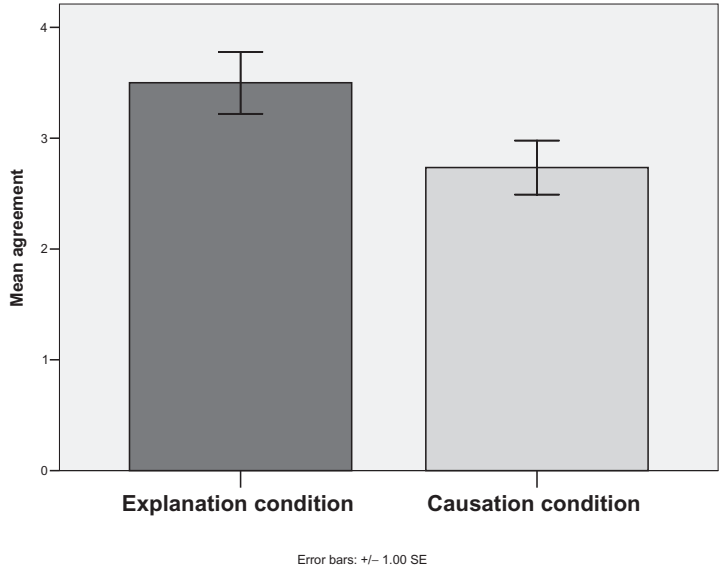

Figure 1. Subjects' mean agreement with the target sentence in Experiment 1. SE, standard error.

to agree with the target sentence in the causation condition than in the explanation condition, not because they thought that absences are not causes, but because they generally prefer sentences formulated with "because" to sentences formulated with "cause."

To control for these confounds, we developed a new case—*the broken rope case*. The broken rope case has the same structure as the rope case, but it involves a genuine event, rather than an absence, in order to control for the possibilities (1) that people simply prefer sentences formulated with "because" to sentences formulated with "cause" and (2) that people are thrown off by the gerundive construction in the causation condition of the rope case.

Experiment 2: The Broken Rope Case
"Susan had to climb a rope in gym class. Susan was a very good climber. She started climbing, but the rope broke before she reached the rafters. She fell on the ground."

In the *causation condition*, this short description was followed by the following question:

"On a scale of 1 to 7, 1 indicating that you totally disagree and 7 indicating that you totally agree, how much do you agree with the following claim? 'The rope breaking caused Susan to fall.'"

In the *explanation condition*, the short description was followed by the following question:

> "On a scale of 1 to 7, 1 indicating that you totally disagree and 7 indicating that you totally agree, how much do you agree with the following claim? 'Susan fell because the rope broke.' "

If the difference between the causation condition and the explanation condition in Experiment 1 were due to the gerundive used in the target sentence of the causation condition ("the rope not breaking") or to a general tendency to agree more with sentences formulated with "because" than with sentences formulated with "cause," we should find a significant difference between the causation condition and the explanation condition of Experiment 2. Since the difference between the causation condition and the explanation condition in Experiment 1 was a medium-size effect, we assumed that a similar effect size would be found in Experiment 2, if one of the two alternative explanations of the finding in Experiment 1 is correct. Assuming such an effect size and an $\alpha = 0.05$, we calculated that we needed a total sample size of $N = 102$ to reach a power of 0.80.

Ninety-five individuals taking classes at the University of Pittsburgh took part in the experiment (mean age: 20; 9 years; range: 18–60 years; 63.2% male).[8] In classroom settings, subjects were randomly assigned to the causation condition ($N = 48$) or to the explanation condition ($N = 47$). Subjects had to answer the relevant question by circling a numeral on a seven-point scale, anchored at 1 with "totally disagree" and at 7 with "totally agree." Subjects were also asked to fill in a short demographic questionnaire. Two subjects answered that they were not native speakers of English. Because excluding these subjects did not affect the data analysis, they were not removed from the data set.

In the causation condition, the mean answer was 5.77 ($SD = 1.37$), while in the explanation condition, the mean answer was 5.47 ($SD = 1.82$). In both conditions, the modal answer was 7. Subjects were not significantly more likely to agree with the target sentence in the explanation condition ("Susan fell because the rope broke") than with the target sentence in the causation condition ("The rope breaking caused Susan to fall") ($t(93) = -0.92$, $p = 0.36$, two-tailed; Figure 2).

Since subjects did not answer differently to the target questions in the causation condition and in the explanation condition of the broken rope case, our findings in Experiment 1 cannot be explained by the two potential confounds: (1) a weaker tendency to agree with sentences formulated with a gerundive and (2) a stronger tendency to agree with sentences formulated with "because."

4.3 The Unsafe Rope Case

We attempted to replicate the findings of Experiment 1. We developed a new case—*the unsafe rope case*. The description of the rope in the rope case does not give any reason to expect the rope to break. By contrast, in the unsafe rope case,

8. Different subjects took part in the experiments described in this article.

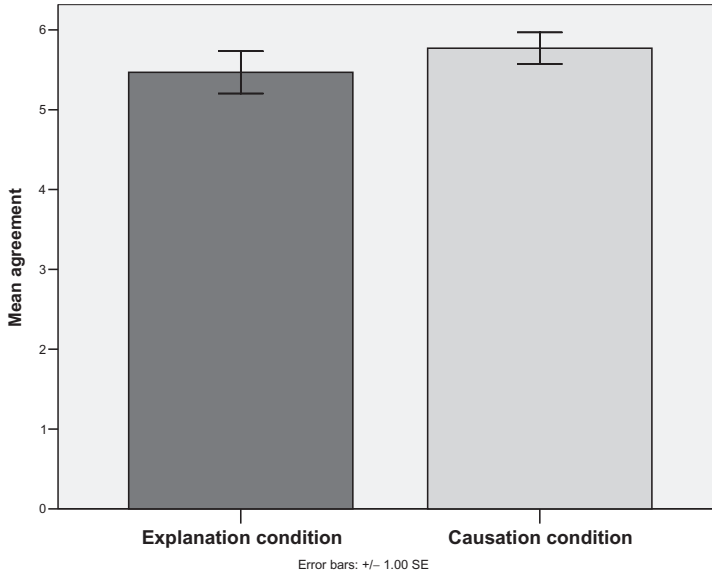

Figure 2. Subjects' mean agreement with the target sentence in Experiment 2. SE, standard error.

the rope is described as being worn out and thus unsafe. Therefore, the rope not breaking has greater explanatory relevance to Susan reaching the rafters in the unsafe rope case than it did in the rope case.

> *Experiment 3: The unsafe rope case*
> "Susan has to climb an old, worn-out rope in gym class. She wondered if it would support her weight. Susan was a very good climber. Though nervous, she climbed all the way to the rafters."

In the *causation condition*, this short description was followed by the following question:

> "On a scale of 1 to 7, 1 indicating that you totally disagree and 7 indicating that you totally agree, how much do you agree with the following claim? 'The rope not breaking caused Susan to reach the rafters.' "

In the *explanation condition*, this short description was followed by the following question:

> "On a scale of 1 to 7, 1 indicating that you totally disagree and 7 indicating that you totally agree, how much do you agree with the following claim? 'Susan reached the rafters because the rope did not break.' "

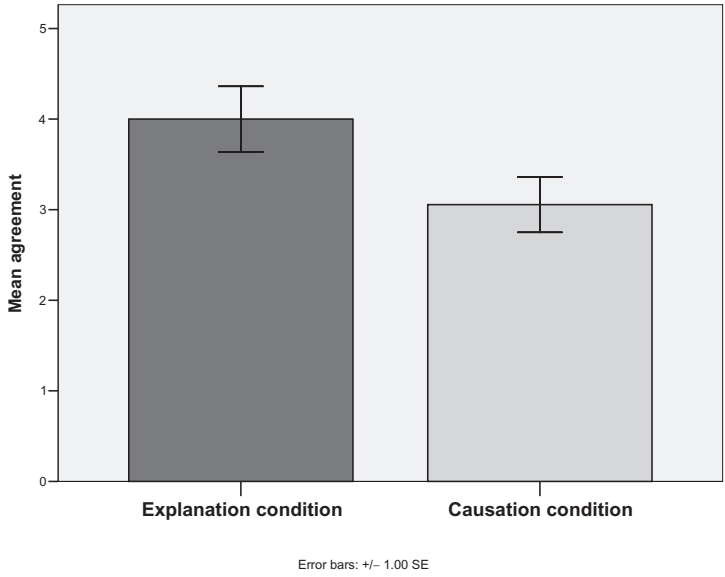

Figure 3. Subjects' mean agreement with the target sentence in Experiment 3. SE, standard error.

Seventy-one individuals taking classes at the University of Pittsburgh took part in the experiment (mean age: 19; 9 years; range: 18–23 years; 58% male). In classroom settings, subjects were randomly assigned to the causation condition (N = 36) or to the explanation condition (N = 35). Subjects had to answer the relevant question by circling a numeral on a seven-point scale, anchored at 1 with "totally disagree" and at 7 with "totally agree." Subjects were also asked to fill in a short demographic questionnaire. Two subjects answered that they were not native speakers of English and one subject failed to answer this question. Because excluding these subjects did not affect the data analysis, they were not removed from the data set.

In the causation condition, the mean answer was 3.06 (SD = 1.84), while in the explanation condition, the mean answer was 4.00 (SD = 2.14). As in Experiment 1, the mean answer in the causation condition was significantly different from a neutral answer, 4 ($t(35) = -3.09$, $p < 0.005$). Subjects were also significantly more likely to agree with the target sentence in the explanation condition ("Susan reached the rafters because the rope did not break") than with the target sentence in the causation condition ("The rope not breaking caused Susan to reach the rafters") ($t(69) = 1.99$, $p = 0.05$, two-tailed; Figure 3). The effect size $d = 0.46$ corresponds approximately to a medium-size effect (Cohen 1992).

4.4 Discussion

Thus, consistent with Beebee's Prediction 1, for some sentences, people are likely to disagree that an absence is a cause, even though the absence qualifies as a

genuine cause according to nonrelationist theories of causation. However, contrary to Beebee's Prediction 2, for some sentences, people are more likely to agree that an event *e* occurs because an absence *a* obtained than they are to agree that *a* caused the occurrence of *e*. This is strong evidence that people do *not* conflate causation and causal explanation.

We should also highlight a striking, unexpected feature of our results. Even when the absence was made clearly explanatorily relevant, as it is in the unsafe rope case, people were not disposed to assent to a sentence causally explaining an event by reference to this absence. In fact, the results of Experiments 1 and 3 are strikingly similar, whereas we intentionally increased the explanatory relevance of the absence in Experiment 3. This might be taken to suggest that people are not disposed to causally explain an event by reference to an absence, merely because this absence is explanatorily relevant. We'll return to this issue in section 7.

5. BEEBEE ON CAUSATION BY ABSENCE

The experiments presented in section 4 point to a significant success as well as a serious challenge for Beebee's solution to the puzzle of causation by absence. We'll consider the former in the next section. We now argue that our experiments suggest a problem that goes to the heart of Beebee's pragmatic solution to the puzzle of causation by absence and leads us to claim that the puzzle remains unsolved. To restate, by supposing both that people confuse causal explanation with causation and that people take inadequate but true explanations to be false, Beebee manages to account for a pattern of behavior (assertion and denial of sentences treating absences as causes) strikingly at odds with her metaphysical theory of causation.

The problem is that our findings suggest that people do not conflate causation and causal explanation. Then, it is not the case, as Beebee claims, that because people conflate explanation and causation, they assert that an absence *a* causes the occurrence of an event *e* whenever they view *a* as explaining the occurrence of *e*—which happens, according to Beebee, when *a* provides a true and conversationally relevant (adequate) explanation of the occurrence of *e*. Furthermore, it is not the case that because people conflate explanation and causation, they deny that an absence *a* causes the occurrence of an event *e* whenever they view *a* as failing to explain the occurrence of *e*—which, according to Beebee, happens when *a* provides a true, but conversationally inappropriate (inadequate) explanation of the occurrence of *e*. Rather, when the folk assert that an absence causes an event, they mean to assert something more than that the absence causally explains the event.[9] Thus, insofar as Beebee's solution requires that the folk *identify* causation with causal explanation, she has failed to solve the puzzle of causation by absence.

But perhaps Beebee does not *need* to commit herself to the claim that the folk identify causation with causal explanation. Perhaps she could commit to something weaker instead. She typically talks as though the proposed confusion between causal explanation and causation amounts to a biconditional: People

9. Like Lewis, Beebee, and other metaphysicians, we assume here that people do sometimes assert that absences are causes. But see section 7 for discussion.

assert that an absence caused an event if and only if they judge the absence to causally explain the event. As we have noted already, this would be inconsistent with our experimental results. However, Beebee might only need the following, weaker claim: People assert that an absence caused an event only if they judge the absence to causally explain the event.

Beebee might even note that our data are perfectly consistent with such a proposal. We found that in the rope case and in the unsafe rope case, people are uncertain with respect to the causal explanation and disagree with the causal sentence. What would be inconsistent with the proposal that people assert that *c* causes the occurrence of *e* only if they judge that *e* occurs because *c* obtains is the inverse combination—people being uncertain with respect to the causal sentence and disagreeing with the causal explanation.

A reply along these lines might look promising, but we do not think that it ultimately succeeds. The relationist cannot merely declare some folk assertions to be false. She also needs to explain *why* the folk make these erroneous assertions. Beebee's original solution tries to do precisely this. While Beebee claims that folk assertions that an absence *a* causes the occurrence of an event *e* are false, she also explains that the folk make these assertions because they conflate causation and causal explanation. And, under her explanation, the folk practice of asserting that some absences are causes comes out as innocuous. The weaker position we are considering on behalf of Beebee (viz. people assert that an absence caused an event only if they judge the absence to causally explain the event) fails to fulfill this requirement. It asserts that the folk make erroneous assertions, without explaining why they do this.

Of course, Beebee can explain folk assertions that some absences are causes by asserting that the folk are just muddle-headed with respect to causation. They have numerous wrong beliefs about causation and a metaphysical theory of causation should steer clear of them. Nothing prevents Beebee from endorsing this line. However, this way of explaining folk judgments of causation by absence would be tantamount to recognizing that relationism largely flies in the face of common sense. Since she appears to judge to be unsatisfying those explanations that leave one's metaphysical position substantially at odds with ordinary judgments, she has not solved the puzzle of causation by absence by her own lights.

6. LEWIS ON CAUSATION BY ABSENCE

Our findings suggest that Beebee's attempt to explain folk judgments of causation by absence fails. Thus, relationism is squarely at odds with part of our commonsense approach to causation. The nonrelationist might therefore conclude that she has ground for rejoicing. We now argue that this is not the case. For, our findings cast doubt on Lewis's explanation of people's failure to mention the vast majority of the causes by absence countenanced by nonrelationist views of causation. Nonrelationism is also at odds with part of our commonsense approach to causation.

Remember that Lewis's crucial move is to note that not saying that *p* is not tantamount to judging that not-*p*. Lewis is certainly right that in general, the former is not identical with the latter. But he might be wrong to distinguish saying that *p*

and judging that not-*p* in the specific case of causation by absence. Beebee had asserted that much, while acknowledging that evidence was lacking to support her assertion (see also McGrath 2005). Our data fill in this gap.

When we (JL and EM) step inside the department of History and Philosophy of Science at the University of Pittsburgh, we do not say that the walls are beige, because this fact is irrelevant. But, suppose now that some undergraduate asks us what the color of the walls is. The question raised by the undergraduate now makes the utterance, "The walls are beige," relevant. We might be surprised by her question, but we would answer that the walls are beige, because we judge that the walls are beige.

Mutatis mutandis, if Lewis's explanation of people's failure to mention certain causes by absence recognized by nonrelationist views of causation is correct, people should agree that these alleged causes are truly causes, when asked. But our experiments suggest that in at least some cases, they don't. Remember that in the causation condition of Experiment 1, we found that people *disagreed* with the sentence

(v) The rope not breaking caused Susan to reach the rafters.

One might reply on behalf of Lewis that people's disagreement with (v) does not mean that they judge that the rope not breaking did not cause Susan to reach the rafters. Rather, by disagreeing with (v), people might mean that the rope not breaking is not the most relevant or important cause of Susan reaching the rafters. This reply should be resisted, however. It is clearly ad hoc: There is no evidence that people take the rope not breaking to be a cause of Susan reaching the rafters at all. Furthermore, as we have seen in Experiment 3, people disagreed with (v) even when it would be relevant to acknowledge a cause by absence. This is prima facie evidence that *pace* Lewis, people's failure to mention numerous causes by absence recognized by nonrelationist views of causation is not merely pragmatic. When people do not say that an absence *a* causes an event *e*—for example, when people do not explain the death of Mick Jagger's plants by asserting that G. W. Bush's failure to water Mick Jagger's plants caused them to die—the reason is often that they judge that *a* did *not* cause *e*—they judge that G. W. Bush's failure to water Mick Jagger's plants did not cause them to die.

If we are right that Lewis's account of folk judgments of causation by absence is incorrect, nonrelationists have to reject part of the commonsense understanding of causation, just as relationists such as Beebee do. Specifically, they have to conclude that the folk practice of judging that absences are causes is partly mistaken in that it denies genuine causes.

7. CONCLUSION: CAUSATION BY ABSENCE

One might have thought that Lewis and Beebee had domesticated the puzzle of causation by absence, albeit in different ways. After all, they have both provided explanations of the folk judgments of causation by absence at odds with their metaphysical views. However, a closer look at folk judgments using experimental methods reveals significant shortcomings in their solutions. *Pace* Lewis, it is not the case that the folk merely fail to utter some sentences involving absences as causes.

Rather, the folk often deny such sentences. *Pace* Beebee, it is not the case that the pattern of folk assertions and denials of sentences that treat absences as causes may be accounted for by supposing that people conflate causation and causal explanation. Rather, the folk do draw this distinction. So, both Lewis's and Beebee's solutions to the puzzle of causation by absence fail to do justice to what the folk think.

Where does this leave us with respect to the debate between relationism and nonrelationism? It leaves us with an unsolved puzzle of causation by absence. Both relationists and nonrelationists have to reject some folk judgments about causation by absence. Relationists have to say that the folk are mistaken when they assert that some absences are causes, and nonrelationists have to say that the folk are mistaken when they say that some absences are not causes, even though these absences have just the right patterns of counterfactual dependence. In order to solve the puzzle, metaphysicians have three choices. First, they may develop yet another account of the pragmatics of ordinary conversation that will account for the folk judgments presented in this article as well as stand up against future research into what the folk think, all while remaining consistent with one or other of the metaphysical choices available. Second, they may develop a competing metaphysical account that respects (or at least more naturally leads to) folk judgments. Or third, they may bite the bullet, admit that the folk are just plain wrong about some of their metaphysical judgments, and provide reasons for accepting some metaphysical account or other *even though* it fails to fit the judgments of ordinary people. We leave it to the metaphysicians to decide which route to take.

By way of conclusion, let us set aside the question of relationism to reflect on the questions that need to be tackled by future research. First, to our surprise, we found that people do not seem disposed to endorse a causal explanation that refers to an absence, merely because this absence is explanatorily relevant. In Experiment 3, people did not agree with the sentence, "Susan reached the rafters because the rope did not break," although the absence—the rope not breaking—had been made explanatorily relevant by describing the rope as worn out and unsafe.

Standard accounts of explanation are not consistent with this finding. According to Lewis's (1986b) account, which is endorsed by Beebee (2004), one truly explains a fact x by means of another fact y if and only if, had it not been the case that y obtained, x would not have obtained either. Such explanations are adequate when the context makes it relevant to mention y. Now, in the unsafe rope case, the fact that the rope did not break is a true causal explanation, since Susan would not have reached the rafters had the rope broke; rather, she would probably have reached the hospital with a broken arm or leg. (Why are philosophical stories always so cruel?) Furthermore, since we emphasized that the rope was worn out, it was relevant to mention the fact that it did not break. By explicitly providing information about the safety of the rope, we made it relevant to explain Susan reaching the rafters by reference to the rope not breaking. Psychologists have emphasized that Grice's (1975) theory of conversational pragmatics applies to experimental contexts that use verbal materials, as we did (e.g., Hertwig and Gigerenzer 1999). When subjects are verbally provided with some information, they typically expect the experimenter to meet some general standards of

communication. Particularly, they expect experimenters to meet what Grice has called the maxim of relevance. That is, the information verbally provided by the experimenter should be relevant to the task at hand. If this is correct, subjects in the explanation condition of the unsafe rope case should have viewed the information about the lack of safety of the rope as relevant to evaluate whether they agreed that Susan reached the rafters because the rope did not break. Thus, in addition to being true, the experimental context made it relevant to assert this explanation. So Lewis's and Beebee's theory of explanation entails that subjects should have agreed with the proposed explanation. But they did not. Thus, to the extent Lewis's and Beebee's account of explanation intends to account for folk explanations, these findings are inconsistent with it. Future research needs to examine this finding in more detail.

Second, we think that partial causation and partial causal explanation ought to be investigated. It might be that the folk deny that some absences are causes when the sentence asserts flatly that some absence causes some event (or that some event happened because an absence obtained). Nevertheless, the folk might affirm that they *are* partial or contributing causes. We did not have time to follow up with versions of our probes that change the target sentence to assert that some absence is a cause of some event (or that some event happened in part because some absence obtained). Such an investigation need not be limited to causation by absence. Indeed, it would be worth knowing under what conditions and to what degree (if at all) people distinguish between causation *simpliciter* and partial causation and between causal explanation *simpliciter* and partial causal explanation.

Third, we think that assumptions everyone takes for granted about how the folk think about causation ought, at last, to be scrutinized. For instance, along with Lewis, Beebee, Mellor, and others, we took for granted that, at least sometimes, the folk assert that some absences are causes, but we did not directly test this claim. After all, one typically does not test what is not in doubt. But having been surprised a few times already, we think it is time to treat our shared assumptions with more systematic care.

Particularly, none of our experiments suggests that the folk take absences to be causes, rather the reverse. Similar results were found in the only relevant study we know of. Knobe (2005) successively gave subjects the two following probes:

> *The Lauren case*
> "Lauren works in the maintenance department of a large factory. It is her responsibility to put oil in the K4 machine on the first day of each month. If she doesn't put in the oil, the machine will break down. On June 1st, Lauren forgot to put in the oil. The machine broke down a few days later."
>
> *The Jane case*
> "Jane also works in the factory, but she does not work in the maintenance department. She works in human resources, keeping track of all the details for the employee health insurance plan. Jane also knew how to put oil in the K4 machine. But no one would have expected her to do so; it clearly wasn't part of her job."

After the Lauren case, subjects were asked to answer the question "Did Lauren *cause* the machine to break down?" on a seven-point scale (0 meaning "no, she didn't" and 6 meaning "yes, she did"). After the Jane case, subjects were asked to answer the question "Did Jane *cause* the machine to break down?" on the same scale. Although the mean answer to the Lauren case was significantly higher than the mean answer to the Jane case, what matters for present purposes is that subjects' mean answer to the Lauren case was only 3.34. That is, the mean answer was just above the middle of the scale, 3. Although Knobe does not comment on this striking result, it appears that in his experiment, just like in ours, people were not clearly disposed to view absences as causal.

So, *is* there a puzzle of causation by absence? *Do* people sometimes judge absences to be causes? Probably. But this has yet to be shown. It might turn out that the folk just don't think absences are causes. We won't know unless we look. The lesson for metaphysicians, as for all of us, is this: The folk probably don't think what you think they think; so rather than guess from the comfort of your armchair, you ought to go out and check.[10]

REFERENCES

Beebee, H. 2004. "Causing and Nothingness." In *Causation and Counterfactuals*, ed. L. A. Paul, E. J. Hall, and J. Collins, 291–308. Cambridge, MA: MIT Press.

Bloom, P. 2004. *Descartes' Baby: How the Science of Child Development Explains What Makes Us Human*. New York: Basic Books.

Cheng, P. W. 1997. "From Covariation to Causation: A Causal Power Theory." *Psychological Review* 104: 367–405.

Cohen, J. 1992. "A Power Primer." *Psychological Bulletin* 112: 155–59.

Davidson, D. 1967. "Causal Relations." *The Journal of Philosophy* 64: 691–703.

Goldvarg, E., and Johnson-Laird, P. N. 2001. "Naïve Causality: A Mental Model Theory of Causal Meaning and Reasoning." *Cognitive Science* 25: 565–610.

Gopnik, A. 1998. "Explanation as Orgasm." *Minds and Machines* 8: 101–18.

Gopnik, A., Glymour, C., Sobel, D. M., Schulz, L. E., Kushnir, T., and Danks, D. 2004. "A Theory of Causal Learning in Children: Causal Maps and Bayes Nets." *Psychological Review* 111: 3–32.

Grice, H. P. 1975. "Logic and Conversation." In *Syntax and Semantics 3: Speech Acts*, ed. P. Cole and J. L. Morgan, 41–58. New York: Academic Press.

Hertwig, R., and Gigerenzer, G. 1999. "The 'Conjunction Fallacy' Revisited: How Intelligent Inferences Look Like Reasoning Errors." *Journal of Behavioral Decision Making* 12: 275–305.

Hirschfeld, L., and Gelman, S., eds. 1994. *Mapping the Mind: Domain Specificity in Cognition and Culture*. New York: Cambridge University Press.

Knobe, J. 2005. "Cognitive Processes Shaped by the Impulse to Blame." *Brooklyn Law Review* 71: 929–37.

Knobe, J., and Fraser, B. Forthcoming. "Causal Judgment and Moral Judgment: Two Experiments." In *Moral Psychology*, ed. Walter Sinnott-Armstrong, Cambridge, MA: MIT Press.

Lewis, D. 1973. "Causation." *The Journal of Philosophy* 70: 556–67.

———. 1986a. *Philosophical Papers*, vol. 2. Oxford: Oxford University Press.

———. 1986b. "Causal Explanation." In *Philosophical Papers*, vol. 2, 214–240. Oxford: Oxford University Press.

10. We would like to thank Josh Knobe and Tania Lombrozo for their comments on this article.

———. 2000. "Causation as Influence." *The Journal of Philosophy* 97: 182–97.
———. 2004. "Void and Object." In *Causation and Counterfactuals*, ed. L. A. Paul, E. J. Hall, and J. Collins, 277–90. Cambridge, MA: MIT Press.
Mandel, D. R. 2005. "Counterfactual and Causal Explanation: From Early Theoretical Views to New Frontiers." In *The Psychology of Counterfactual Thinking*, ed. D. R. Mandel, D. J, Hilton, and P. Catellani, 11–23. New York: Routledge.
McGrath, S. 2005. "Causation by Omission." *Philosophical Studies* 123: 125–48.
Mellor, D. H. 1995. *The Facts of Causation*. London: Routledge.
———. 2004. "For Facts as Causes and Effects." In *Causation and Counterfactuals*, ed. L. A. Paul, E. J. Hall, and J. Collins, 309–23. Cambridge, MA: MIT Press.
Nahmias, E. A., Morris, S. G., Nadelhoffer, T., and Turner, J. 2006. "Is Incompatibilism Intuitive?" *Philosophy and Phenomenological Research* 73: 28–53.
Nichols, S. 2004. "Folk Concepts and Intuitions: From Philosophy to Cognitive Science." *Trends in Cognitive Sciences* 8: 514–18.
Peirce, C. S. 1994. *The Collected Papers of Charles Sanders Peirce*. Charlottesville, VA: InteLex Corporation. Available at http://library.nlx.com/titles/index.cfm.
Rips, L., Blok, S. V., and Newman, G. 2006. "Tracing the Identity of Objects." *Psychological Review* 113: 1–30.
Salmon, W. C. 1984. *Scientific Explanation and the Causal Structure of the World*. Princeton, NJ: Princeton University Press.
Sloman, S. 2005. *Causal Models: How People Think about the World and Its Alternatives*. Oxford: Oxford University Press.
Spellman, B. A., Kincannon, A., and Stose, S. 2005. "The Relation between Counterfactual and Causal Reasoning." In *The Psychology of Counterfactual Thinking*, ed. D. R. Mandel, D. J. Hilton, and P. Catellani, 28–43. London: Routledge Research.
Wolff, P. 2007. "Representing Causation." *Journal of Experimental Psychology: General* 136: 82–111.