King's Research Portal

DOI:

[10.1080/09515089.2013.815099](https://doi.org/10.1080/09515089.2013.815099)

Document Version

Peer reviewed version

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Clarke, R., Shepherd, J., Stigall, J., Waller, R. R., & Zarpentine, C. (2015). Causation, Norms, and Omissions: A Study of Causal Judgments. *PHILOSOPHICAL PSYCHOLOGY*, 28(2), 279-293.
<https://doi.org/10.1080/09515089.2013.815099>

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Causation, Norms, and Omissions: A Study of Causal Judgments

Randolph Clarke, Joshua Shepherd, John Stigall, Robyn Repko Waller, and Chris Zarpentine

Abstract: Many philosophical theories of causation are egalitarian, rejecting a distinction between causes and mere causal conditions. We sought to determine the extent to which people's causal judgments discriminate, selecting as causes counternormal events—those that violate norms of some kind—while rejecting non-violators. We found significant selectivity of this sort. Moreover, priming that encouraged more egalitarian judgments had little effect on subjects. We also found that omissions are as likely as actions to be judged as causes, and that counternormative selectivity appears to apply equally to actions and omissions.

Philosophical discussion of causation often exhibits a notably egalitarian spirit. It is widely held that for a given outcome there are typically many causes (at each of the prior times at which there's any cause). Some causes might be more salient than others, more likely to be noticed or cited (when we do notice such things) as causes or *the* causes, while others might be regarded as mere conditions. But, it is maintained, these are pragmatic matters. As far as metaphysics is concerned, all are equally causes.

David Lewis provides a crisp example of this approach:

We sometimes single out one among all the causes of some event and call it “the” cause, as if there were no others. Or we single out a few of the “causes,” calling the rest mere “causal factors” or “causal conditions.” Or we speak of the “decisive” or “real” or “principle” cause. We may select the abnormal or extraordinary causes, or those under human control, or those we deem good or bad, or just those we want to talk about. I have nothing to say about these

principles of invidious discrimination. I am concerned with the prior question of what it is to be one of the causes (unselectively speaking). (Lewis, 1973, pp. 558-559)

Thus, while we might ordinarily not say that the presence of oxygen caused the house fire—and perhaps not even that it was *a* cause of the fire—on the view in question it is as much a cause as is the act of arson, the dropped cigarette, or the lightning strike.

It is often observed that among the factors that influence which among the many causes of an outcome might be mentioned are norms of various kinds and what is, in one way or another, normal (Hart and Honoré, 1959). Some things that happen are not supposed to happen, or they usually do not. Let's call something of this sort *counternormal*. In our selection of causes, we often pick out something that was counternormal. There is always oxygen in and around the house. There is not supposed to be an act of arson, or a dropped cigarette, and lightning strikes are unusual.

A few philosophers have argued, against this trend, that norms or what is normal play a role not just in our selection among causes but in what is a cause in the first place. Sometimes this view is urged only in the case of causation involving absences or omissions. Consider:

Barry promises to water Alice's plant, doesn't water it, and the plant then dries up and dies. Other people who have never heard of Alice or her plant did not water it either. Carlos (who lives far away in Australia) did not water it, Dirk did not water it, Eric did not water it, and so on. (McGrath, 2005, p. 126)

In such a case, it is argued, it might well be that Barry's not watering the plant caused its death, while the other failures to water the plant weren't causes of that outcome. The difference is said to hinge on the fact that of the several people who didn't water the plant, only Barry was supposed to water it.

Sometimes it is argued, more broadly, that norms or what is normal figure *generally* in determining which causally relevant things were causes of some outcome. On this view, our concept of causation is inequalitarian. Hitchcock and Knobe (2009) find support for this view in studies in which participants responded to the following scenario:

The receptionist in the philosophy department keeps her desk stocked with pens.

The administrative assistants are allowed to take pens, but faculty members are supposed to buy their own.

The administrative assistants typically do take the pens. Unfortunately, so do the faculty members. The receptionist repeatedly e-mailed them reminders that only administrators are allowed to take the pens.

On Monday morning, one of the administrative assistants encounters Professor Smith walking past the receptionist's desk. Both take pens. Later that day, the receptionist needs to take an important message...but she has a problem. There are no pens left on her desk. (p. 594)

In this case, if Professor Smith hadn't taken a pen, there wouldn't have been a problem; and if the administrative assistant hadn't taken a pen, there wouldn't have been a problem. Even so, when presented with statements concerning who caused the problem, participants tended to agree that Professor Smith caused the problem ($M = 5.2$ on a 1 to 7 scale) and disagree with the statement that the administrative assistant caused the problem ($M = 2.8$).

Hitchcock and Knobe were able to replicate these results in a case not focused on intentional agency:

A machine is set up in such a way that it will short circuit if both the black wire and the red wire touch the battery at the same time. The machine will not short circuit if just one of these wires touches the battery. The black wire is designated as the one that is supposed to touch the battery, while the red wire is supposed to remain in some other part of the machine.

One day, the black wire and the red wire both end up touching the battery at the same time. There is a short circuit. (2009, p. 604)

In this case, if the black wire hadn't touched the battery, there wouldn't have been a short circuit; and if the red wire hadn't touched the battery, there wouldn't have been a short circuit. Even so, participants tended to agree that the fact that the red wire touched the battery caused the machine to short circuit ($M = 4.9$) and disagree with the statement that the fact that the black wire touched the battery caused the machine to short circuit ($M = 2.7$).

It is a matter of some controversy whether absences or omissions are causes (or effects) at all (Armstrong, 2004; Beebe, 2004; Dowe, 2001; Lewis, 2004b; Schaffer, 2000). And some experimental testing of causal judgments has found relatively low acceptance of statements identifying absences as causes. Livengood and Machery (2007) had participants respond to several scenarios; we describe two of them here.

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. (p. 115)

Experiment 3: The unsafe rope case

Susan has [sic] 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. (p. 119)

Most of the participants in both experiments disagreed with the claim that the rope not breaking caused Susan to reach the rafters. ($M = 2.73$ on a scale of 1 to 7 in Experiment 1; $M = 3.06$ in Experiment 3.) There was no significant difference between Experiment 1 and Experiment 3 in participants' responses to this statement, despite the fact that the scenario in Experiment 3 suggested that indeed the rope might break.

We take no stand here on whether norms or what is normal figure in distinguishing between causes and mere conditions, that is, on whether causation is egalitarian or not. However, we wanted to see whether counternormativity plays a significant role in people's causal

judgments, whether it plays a larger role in judgments about omissions than it does in judgments about actions, and whether participants might be primed in such a way that the role of counternormativity was reduced. The results we obtained were surprising in several respects.

While people's causal judgments in some manner reflect their concept(s) of causation, caution is in order if conclusions about concepts are to be drawn from studies such as ours. We discuss this issue in the final section of the paper.

1. First Experiment

There are ways of phrasing questions or statements about causes that invite selection. Participants can be presented with a statement that such-and-such was *the* cause of a certain outcome, which of course encourages discrimination. Even a statement that such-and-such *caused* the outcome might encourage selectivity. We wanted to see what participants would say when the prompt didn't invite selection. Hence our target statements always said that such-and-such was *one of the causes* of the outcome. If people employ an egalitarian concept of causation, the chosen phrasing would seem suited to elicit it.¹

Further, some of our participants received additional information that encouraged egalitarianism. We wanted to see if this additional information influenced agreement with our target statements.

In our first study, all of the participants read the following vignette about a collision:

¹ Compare Livengood and Machery: "It might be that the folk deny that some absences are causes when the sentence asserts flatly that some absence causes some event.... Nevertheless, the folk might affirm that they *are* partial or contributing causes" (2007, p. 125). They go on to suggest that the effect might not be limited to absences, but might appear also when an event is said to cause or be the cause.

Two cars, one driven by Greta and the other driven by Rachel, were approaching an intersection. Greta had a green light. Rachel had a red light, but she wasn't paying attention. The lights stayed that way. Neither driver stopped, and their cars collided.

A third of the participants also read a short prime (immediately following the vignette) suggestive of an egalitarian view about the causes of the collision ("Many things were causes of the collision"). Another third received a more extensive prime that included a counterfactual claim about the collision (e.g., "Many things were causes of the collision. And the collision wouldn't have occurred if things had differed in certain ways. For example, it wouldn't have occurred if Greta hadn't driven into the intersection"). The action or omission mentioned in the counterfactual claim was the same as the action or omission mentioned in the statement that the participant received (see the statements below). The remaining third of the participants were not primed.

Participants then read exactly one of the following four target statements about causes of the collision and indicated their agreement on a seven-point Likert scale (from 1 as "strongly disagree" to 7 as "strongly agree," with 4 as "neutral"):

Rachel's driving into the intersection was one of the causes of the collision.

Rachel's not stopping was one of the causes of the collision.

Greta's driving into the intersection was one of the causes of the collision.

Greta's not stopping was one of the causes of the collision.

We predicted that if we primed participants with information suggestive of egalitarianism (that the collision had many causes, and that what Greta did made a difference to whether the collision occurred), then participants' responses would be more reflective of egalitarianism. Thus, we expected to find more agreement among the primed participants that Greta's driving into the intersection or Greta's not stopping was one of the causes of the collision, despite the fact that Greta had a green light. We also hypothesized that participants would tend to favor actions (e.g., Rachel's driving into the intersection) over omissions (Rachel's not stopping) as causes of the collision, and we expected that priming might decrease the extent of this favoring.

The study utilized a 2 (Norms: Rachel vs. Greta) by 2 (Description: action vs. omission) by 3 (Prime: no prime vs. short prime vs. extensive prime) design. Participants in the study were 587 workers on Amazon's Mechanical Turk, with roughly 50 participants per condition. We ran a 2x2x3 between-subjects analysis of variance (ANOVA) for the level of agreement with the statements about causes of the collision.

There was a main effect of Norms ($F(11, 575) = 297.437, p < .001$). That is, participants more strongly agreed that Rachel was one of the causes of the collision than that Greta was one of the causes (see Fig. 1). The mean for Rachel as one of the causes overall was 5.96 ($SD = 1.55$), whereas the mean for Greta overall was 3.30 ($SD = 2.146$). There were no main effects of Prime ($F(11, 575) = .048, p = .953$) or Description ($F(11, 575) = 2.123, p = .146$), and there were no interactions. Hence, even when primed, participants tended to judge Rachel to be among the causes and Greta not to be; the causal judgments were stubbornly inegalitarian. Moreover, actions and omissions were judged the same with regard to their causal status. Agreement that Rachel's not stopping was one of the causes was not significantly different from agreement that Rachel's driving into the intersection was one of the causes, and likewise for Greta.

[Figure 1 here]

These surprising results suggest several features of people's causal judgments. First, it appears that norms have a significant influence on these judgments. Rachel violated the norm that one must stop at a red light, and Greta did not. No matter how plain we made it that many things were causes of the outcome or that what Greta did made a difference to whether the outcome occurred, participants tended to take Rachel but not Greta to be one of the causes. Norm violation seemed to sway causal judgments even in the face of information suggestive of egalitarianism. This result is consistent with Hitchcock and Knobe's (2009) view that people have an inequalitarian concept of causation. Moreover, our results go beyond Hitchcock and Knobe's findings to support the view that people are surprisingly resistant, even given extensive suggestion, to agreeing that agents who haven't violated norms might nevertheless be among the causes of some outcome.

It should be noted, however, that although there was a highly significant main effect of Norms in designating causes of the collision, there was a sizeable subset of the participants (about 38%) that agreed at least slightly that Greta was one of the causes of the collision. This result is consistent with at least a subgroup of the population employing an egalitarian notion of causation.

It is possible that, although all participants possessed an egalitarian concept, many misapplied that concept, or (despite the wording of the target statements) applied a different concept, perhaps led to do so by pragmatic factors. The norms in play were moral and legal, and the scenario strongly invited judgments of blame. Attributions of causation and of blame are

closely associated. Indeed, attributions of causation can implicate blame. Some participants might have rejected statements that Greta was one of the causes of the collision because of the pragmatic infelicity or inadequacy of such statements, and despite the fact that their concept of causation applied to Greta in this case. Although we think this possibility worth considering, our study did not provide a way of testing it. We recommend it for further study.²

Second, the lack of a main effect for Description suggests several competing hypotheses about the status of omissions as causes. It could be that people see omissions, such as an agent's not stopping at a light, to be on equal footing with actions as candidate causes of outcomes. Another possibility consistent with our results is that although people more readily identify actions as causes than they do omissions when norms aren't involved, they draw no such distinction when norms are salient. Since Rachel was legally required to stop, it appears not to have mattered to participants whether her conduct was characterized as "driving into the intersection" or as "not stopping"; and likewise, since Greta was not legally at fault for the collision, it seemed not to matter whether her conduct was characterized as an action or as an omission. A third possibility in line with the lack of a main effect of Description is that because an agent's action of driving into an intersection is so easily redescribed as her not stopping, any distinction between actions and omissions was obscured.

² Hitchcock & Knobe argue that "people's causal intuitions can be affected by norm violations, even in the absence of any judgment of blameworthiness" (2009, p. 604). Their short circuit case is offered in support of this claim. They take a judgment of blameworthiness to be a moral judgment. But the notion of *being to blame* is broader and is arguably applicable in the short circuit case. They also found that counternormative events were favored as causes even when the outcome was a happy one, and hence when nothing was to blame for it. Any successful defense of egalitarianism of the sort we consider in the text here would need to be extended to explain responses in these cases as well.

2. *Second Experiment*

To test which of the above explanations of the lack of a main effect of Description is correct, we ran a second study. Here we wanted to test whether clarifying the difference between an action and an omission would yield a main effect of Description. We modified the collision story slightly so that the action was not redescribed so easily as an omission.

Action condition

Two cars, one driven by Greta and the other driven by Rachel, were approaching an intersection. Greta had a green light. Rachel had a red light, but she wasn't paying attention and turned left in the intersection. Unaware of Rachel, Greta did not step on her brakes. The lights stayed that way, and the cars collided in the intersection.

After reading this vignette participants read one of two statements about causes of the collision and indicated their agreement on a seven-point Likert scale (again, from 1 as "strongly disagree" to 7 as "strongly agree," with 4 as "neutral"):

Greta's not stepping on her brakes was one of the causes of the collision.

Rachel's turning left in the intersection was one of the causes of the collision.

To get a clear contrast between counternormative action and omission, we presented some participants with the following:

Omission condition

Two cars, one driven by Greta and the other driven by Rachel, were approaching an intersection. Greta had a green light. Rachel had a red light, but she wasn't paying attention. The lights stayed that way. Unaware of each other, neither driver stepped on the brakes. Both cars collided.

After reading this vignette participants read one of two statements about causes of the collision and indicated their agreement on a seven-point Likert scale:

Greta's not stepping on her brakes was one of the causes of the collision.

Rachel's not stepping on her brakes was one of the causes of the collision.

We ran a 2 (Norms: Rachel or Greta) by 2 (Description: Action or Omission) between-subjects ANOVA on agreement with the statements about causes of the collision. The participants were 197 workers on Amazon's Mechanical Turk, with roughly 50 participants per condition.

The results replicated the main effect of Norms from the first study ($F(3, 193) = 76.22, p < .001$). The overall mean for Rachel as one of the causes of the collision was 6.03 ($SD = 1.79$), whereas the overall mean for Greta as one of the causes was 3.48 ($SD = 2.26$) (see Fig. 2). This result further supports the view that normativity plays an important role in people's causal judgments.

[Figure 2 here]

However, there was no main effect of Description ($F(3, 193) = .202, p = .654$).

Participants were no more likely to count an agent's conduct as a cause of the collision when it was described as an action of turning left as opposed to an omission to step on the brakes (e.g., mean of 6.08 ($SD = 1.58$) for Rachel's turning left versus mean of 5.98 ($SD = 2.00$) for Rachel's not stepping on her brakes). This suggests that people will identify omissions as causes just as readily as they will identify actions as causes. Whereas participants in Livengood and Machery's (2007) study did not identify the rope's not breaking—an absence—as a cause of Susan's climbing to the top even when this absence was made salient in the vignette, our study provides evidence that people will pick out some omissions as causes, at least when there is a relevant norm involved. Of course, there is an important difference between an absence of a non-agentive event and an omission—viz., the agentive aspect of the latter. Thus, it is open for future studies to explore whether the agentive nature of omissions is a key feature that leads people to count some omissions as causes but not to so count other absences. However, it seems plausible that people recognize at least some absences as causes—particularly counternormal ones, such as a severe lack of rainfall as a cause of famine or a complete lack of food as a cause of death.

3. Third Experiment

There is a remaining alternative explanation of the lack of a main effect of Description in study two, viz., the possibility explored earlier that people discriminate against omissions as causes only in scenarios that do not involve salient norms. To see whether this explanation might be ruled out, we ran a third study in which neither driver was legally at fault for the collision. Participants in this third study read the following vignette:

Two cars, one driven by Greta and the other driven by Rachel, were approaching an intersection. The traffic light had malfunctioned, and both drivers had a green light. The light stayed that way, and the cars collided in the intersection.

Participants then read exactly one of the following statements about causes of the collision and indicated their agreement on a seven-point Likert scale:

Rachel's driving into the intersection was one of the causes of the collision.

Rachel's not stopping was one of the causes of the collision.

The malfunctioning of the light was one of the causes of the collision.

Note that as in the vignette from study one, both Rachel and Greta drive into the intersection in this case. However, in contrast to the vignettes from both study one and study two, neither driver drives through a red light; rather, a light malfunction results in both drivers having a green light. This type of scenario, then, should tease apart the influence of norms from any influence an action/omission distinction might have. If, for example, participants evidence an action/omission distinction when neither driver is legally at fault, this would support the claim that norms are far more important for people's causal judgments than the (still drawn) action/omission distinction is. On the other hand, if there is no effect of the action vs. omission description, as in our previous studies, this would support the view that the action/omission distinction is not highly relevant for causal judgments, in a broader group of cases than simply those involving norm violation.

We ran a one-way between-subjects ANOVA for level of agreement with the statement about causes of the collision with Description (Rachel's action or Rachel's omission or the

light's malfunction) as the factor. Our participants were 144 workers on Amazon's Mechanical Turk, with roughly 50 participants per condition. There was a significant effect of Description ($F(2, 141) = 12.297, p < .001$). The mean rating of agreement for Rachel's driving into the intersection—Rachel's action—as one of the causes of the collision was 4.24 ($SD = 2.087$). This was comparable to the mean for Rachel's not stopping—Rachel's omission—as one of the causes of the collision, which was 4.58 ($SD = 2.239$). In contrast, participants gave the highest level of agreement to the statement that the light's malfunction was one of the causes of the collision ($M = 6.13, SD = 1.513$) (see Fig. 3). Post-hoc pairwise comparisons using the Tukey HSD test confirm that the mean rating of agreement with the light's malfunction as one of the causes of the collision was significantly different from both the mean rating of agreement with Rachel's driving into the intersection as one of the causes ($p < .001$) and the mean of Rachel's not stopping as one of the causes ($p = .001$). However, the mean rating of agreement with Rachel's driving into the intersection as one of the causes of the collision was not significantly different from the mean for Rachel's not stopping as one of the causes. Hence, even when certain norms, such as legal norms of fault, are absent, participants do not distinguish between actions and omissions as candidate causes of events. These results further support the claim that people take actions and omissions to be on equal footing when it comes to designating causes of an outcome.

[Figure 3 here]

Moreover, the fact that participants tended more to agree that the light's malfunction was one of the causes of the collision than to agree that Rachel's action or omission was one of the

causes indicates that norms beyond agentic norms, such as the proper functioning of some artifact, feature as a consideration when designating causes of an outcome. This finding is in line with Hitchcock and Knobe's (2009) results reported earlier.

4. Fourth Experiment

Thus far, our studies suggest that counternormativity has a significant influence on people's causal judgments, but that actions and omissions have equal status as candidates for causes. Although priming participants toward a more egalitarian notion of causation—one that considers both normal and counternormal things as among the causes of an outcome—failed to elicit egalitarian causal judgments in our first study, we wanted to make sure that the lack of effect for the prime was not a function of the particular way in which we presented the egalitarian suggestions. Hence we ran a further study that employed another method of encouraging participants to make egalitarian judgments. Participants were 81 workers on Amazon's Mechanical Turk, with roughly 40 workers per condition. Each participant read the same story as the no-prime condition in study one, but with one of the following additions:

Greta condition

A physics student read about the collision in the newspaper. The student said:

‘Of course, Rachel is legally to blame for the collision.’ The student then added:

‘Still, considering the physics of the situation, Greta's driving into the intersection was also one of the causes of the collision.’

Rachel condition

A physics student read about the collision in the newspaper. The student said:

‘Of course, Rachel is legally to blame for the collision.’ The student then added:

‘Moreover, considering the physics of the situation, Rachel’s driving into the intersection was one of the causes of the collision.’

Participants were then asked to indicate their agreement on a 7-point Likert scale with one of the following statements:

Greta condition

Considering the physics of the situation, Greta’s driving into the intersection was one of the causes of the collision.

Rachel condition

Considering the physics of the situation, Rachel’s driving into the intersection was one of the causes of the collision.

The additions to the original vignettes and the added clause to the statements about the causes of the collision frame the issue of causation from a physics perspective. We reasoned that when participants turn from issues about who is legally at fault and focus instead on the physics of the collision, they might be more likely to give egalitarian responses—to agree that Greta’s driving into the intersection was one of the causes of the collision, despite the fact that Greta had a green light. However, if participants were still reluctant to agree that Greta’s action was one of the causes of the collision, this result would strengthen support for the claim that counternormativity

plays a robust role in people's causal judgments.

An independent samples t-test was conducted on the effect of Norms (Rachel or Greta) on level of agreement with the statements presented to participants. The mean rating of agreement with Greta's driving into the intersection as one of the causes of the collision ($M = 4.62$, $SD = 1.858$) was significantly different from the mean for Rachel's driving into the intersection as one of the causes ($M = 5.71$, $SD = 1.672$; $t [79] = -2.802$, $p = <.01$). Hence, it appears that people are still resistant, even when invited to assess a scenario from a physics perspective, to treating normal events as on par with counternormal events as candidate causes (see Fig. 4). However, the addition of the physics prime did have an effect. Participants in our first experiment averaged a 3.30 agreement rating with statements that feature Greta as one of the causes of the collision—below the midpoint “neutral” response, and in the disagreement range. In contrast, participants in our fourth experiment averaged a 4.62 agreement rating with the statement that Greta was one of the causes of the collision—above the midpoint, and in the agreement range. Thus, it seems that once participants were given a justification from physics for designating normal events as causes, they were more willing to do so.

[Figure 4 here]

5. Discussion

Our experiments confirm and expand on previous results indicating that norms or what is normal influence causal judgments. In fact, people seem highly resistant to egalitarianism. Even when egalitarianism is strongly suggested by prompts, participants' judgments appear to be dominated

by counternormativity.

At the same time, some participants in all of our experiments agreed that normal factors were causes. And in our last experiment we found that focusing participants' attention in a certain way—emphasizing the physics of a situation—can increase the frequency with which they give egalitarian responses.

In our third experiment, participants agreed to a far greater extent that the traffic light malfunction was a cause than either that Rachel's action or that her omission was a cause. Consistent with the results reported by Hitchcock and Knobe (2009), this result indicates that the normativity relevant to causal judgments is not limited to agentic norms. It remains to be determined how different kinds of norms might interact and whether agentic norms influence causal judgments to a greater degree than other kinds of norms (e.g., statistical norms or norms of proper function).

We also found no evidence that whether someone's conduct is characterized as an action or as an omission affects causal judgment. However, further testing is needed to provide more definitive support here. It is possible that the action and omission descriptions used in our studies were not sufficiently different to lead participants to represent them differently; perhaps a more dramatic difference between action and omission might lead people to judge such cases differently.

While Hitchcock and Knobe did find evidence that norms matter even in cases that involve no agents, it remains to be seen whether agentic omissions and absences of other sorts are treated the same. Mere salience did not lead people to treat an absence as a cause in Livengood and Machery's (2007) study. However, it is likely that a counternormative absence would be judged a cause just as a counternormative omission was judged to be a cause in our

study.

We began our paper with a dispute concerning the nature of causation—as it might be said, a question of metaphysics. Our studies report people’s causal judgments concerning various scenarios. One might think that the studies reveal something about people’s concept of causation, but that the latter is quite independent from the metaphysical question. We find both of these issues rather more complicated.

The egalitarian we quoted, Lewis, makes it clear that he takes himself to be engaged in conceptual analysis.³ He advances his theory in answer to the question: “What is causation? As a matter of analytic necessity, across all possible worlds, what is the unified necessary and sufficient condition for causation?” (2004b, p. 287). The question presumes a tight connection between conceptual analysis and metaphysics.

One view of concepts in line with this presumption has it that, with the exception of primitives, each concept has a structure that provides necessary and sufficient conditions for its applicability (where a concept of *F* is applicable to all and only *F*s). On this view, correct analysis of a concept yields analytic truths about the concept’s referent (if it has one). And, of course, analytic truths (if there are any) have clear implications for metaphysics. If it is an analytic truth that all *F*s are *G*, then it is a truth—and a metaphysically necessary one—that all *F*s are *G*.

On such a view, there’s no room for incompatible concepts—concepts with incompatible structures—of a given phenomenon. Concepts of a single thing that had incompatible structures would yield inconsistent analyses. And there can’t be inconsistent truths, much less inconsistent

³ See especially his 2004a, pp. 75-79 and 2004b, pp. 287-289.

analytic truths. Given this view, if Lewis's analysis is correct, no one has an inegalitarian concept of *causation* (though people might have inegalitarian concepts of related phenomena—salient causation, for example). On the other hand, if people's concept of causation is inegalitarian, then, on this view of concepts, Lewis is mistaken about the nature of causation.

Other conceptions of concepts do allow for structurally incompatible concepts of a given phenomenon. On one such view, a concept's structure is its role in someone's theory, something determined holistically by its place within a network. As an individual's theory of things changes, so will the structures of her concepts change; and individuals with different theories can have structurally different concepts of some phenomenon. Since this view allows that various concepts of a given phenomenon can have mutually incompatible structures, it severs the connection between conceptual structure and truth: we cannot safely move from the structure of someone's concept of *F* to the nature of *F*s.⁴

If we ask what studies such as ours reveal about people's concepts, we should be clear about which view of concepts we're employing in our inquiry.

Given the first view, our studies present a challenge to egalitarians. What people say about cases is by no means irrelevant to the project of conceptual analysis; indeed, people's judgments provide one kind of datum for this endeavor. However, it is not obvious that egalitarians can't meet the challenge.

As we observed, a significant number of participants responded in ways consistent with

⁴ For more on these (and other) views of concepts, see Laurence & Margolis 1999. In fact Lewis's own view of conceptual analysis allows for variation in concepts of a given phenomenon, and he allows that the referent of a given concept might be an imperfect (but good enough) fit for the analysis of that concept. See, for example, his 1989, pp. 130-131; and for discussion of his views on these matters, see Nolan 2005, pp. 213-228. We describe a more classical view of concepts and conceptual analysis in order to present a sharp distinction between competing views of concepts.

egalitarianism, and our final study found a priming effect suggestive of egalitarianism. True enough, a full defense of egalitarianism will need to explain the widespread rejection of causal claims that, given egalitarianism, are true. We suggested one thing to which such an explanation might appeal—a conflation of pragmatic infelicity with falsehood. And we noted difficulties that an effort of this sort faces: the absence of a priming effect in our first study and the weakness of the effect in study four.

If we employ the second view of concepts, the inquiry concerning implications of our studies will be rather different. People’s causal judgments provide strong evidence concerning the structures of their concepts, on this view. But our inquiry will then have no clear bearing on Lewis’s proposal, for its results won’t bear directly on the nature of causation.

Drawing conclusions about people’s concept(s) of causation from studies such as ours is thus complicated by the fact that there are rather different conceptions of concepts.⁵ We do not (as a group) endorse any particular conception. It is for this reason that we have for the most part couched discussion of our studies in terms of people’s *judgments*, rather than in terms of their *concepts*. The further questions are, we agree, both interesting and important. But it will take more in the way of both empirical work and philosophical theory to settle them. We recommend such further inquiry, with attention to the conception(s) of concepts employed, and offer our results as data to be taken into account in its conduct.

⁵ Alexander, Mallon, and Weinberg (2010) present this and other difficulties for what they call the “positive program” of experimental philosophy, the effort to answer traditional philosophical questions by means of experimental methods such as the administration of surveys.

References

- Alexander, J., Mallon, R., & Weinberg, J. (2010). Accentuate the negative. *Review of Philosophy and Psychology*, 1, 297-314.
- Armstrong, D. M. (2004). Going through the open door again. In J. Collins, N. Hall, & L. A. Paul (Eds.), *Causation and counterfactuals* (pp. 445-457). Cambridge, MA: MIT Press.
- Beebe, H. (2004). Causing and nothingness. In J. Collins, N. Hall, & L. A. Paul (Eds.), *Causation and counterfactuals* (pp. 291-308). Cambridge, MA: MIT Press.
- Dowe, P. (2001). A counterfactual theory of prevention and “causation” by omission. *Australasian Journal of Philosophy*, 79, 216-226.
- Hart, H. L. A., & Honoré, A. M. (1959). *Causation in the law*. Oxford: Clarendon Press.
- Hitchcock, C., & Knobe, J. (2009). Cause and norm. *Journal of Philosophy*, 106, 587-612.
- Laurence, S., & Margolis, E. (1999). Concepts and cognitive science. In E. Margolis & S. Laurence (Eds.), *Concepts: Core readings* (pp. 3-81). Cambridge, MA: MIT Press.
- Lewis, D. K. (1973). Causation. *Journal of Philosophy*, 70, 556-567.
- Lewis, D. K. (1989). Dispositional theories of value. *Proceedings of the Aristotelian Society, Supplementary Volumes*, 63, 113-137.
- Lewis, D. K. (2004a). Causation as influence. In J. Collins, N. Hall, & L. A. Paul (Eds.), *Causation and counterfactuals* (pp. 75-106). Cambridge, MA: MIT Press.
- Lewis, D. K. (2004b). Void and object. In J. Collins, N. Hall, & L. A. Paul (Eds.), *Causation and counterfactuals* (pp. 277-290). Cambridge, MA: MIT Press.
- Livengood, J., & Machery, E. (2007). The folk probably don’t think what you think they think: Experiments on causation by absence. *Midwest Studies in Philosophy*, 31, 107-127.

McGrath, S. (2005). Causation by omission: A dilemma. *Philosophical Studies*, 123, 125-148.

Nolan, D. 2005. *David Lewis*. Chesham: Acumen.

Schaffer, J. (2000). Causation by disconnection. *Philosophy of Science*, 67, 285-300.

Figures

Figure 1. First Experiment, means for Rachel (norm violation) and Greta (no norm violation) by Description condition: No prime, Short prime, and Extensive prime.

Figure 2. Second Experiment, means for Rachel (norm violation) and Greta (no norm violation) by condition: Action and Omission.

Figure 3. Third Experiment, means by condition: Action, Omission, or Norm violation by a traffic light.

Figure 4. Fourth Experiment, means for Greta and Rachel.