Highlighting the Causal Meaning of Causal Test Questions in Contexts of Norm Violations

Jana Samland (jana.samland@psych.uni-goettingen.de)
Michael R. Waldmann (michael.waldmann@bio.uni-goettingen.de)
Department of Psychology, University of Göttingen, Gosslerstr. 14, 37073 Göttingen, Germany

Abstract

Experiments have shown that prescriptive norms often influence causal inferences. The reason for this effect is still not clear. One problem of the studies is that the term 'cause' in the test questions is ambiguous and can refer to both the causal mechanism and the agent's accountability. Possibly subjects interpreted the causal test question as a request to assess accountability rather than causality. Scenarios that put more stress on the causal mechanism should therefore yield no norm effect. Consequently, Experiment 1 demonstrates that norms no longer influence causal judgments when the causal information is presented in a trial-by-trial learning task. Furthermore, Experiment 2 shows that norm effects are only obtained when the test question asks about a (potentially accountable) person but not when asked about a component of the causal mechanism. Both findings demonstrate that norms cease to influence causal judgments when the task settings highlight causal relations.

Keywords: causal reasoning; moral judgment; causal selection; norms

Introduction

Recent findings in the philosophical and psychological literature have challenged the traditional view about the relationship between causality and morality: it is not only the case that causal inferences influence moral evaluations but also, in reverse, that causal inferences are influenced by moral assessments (Alicke, 1992; Alicke, Rose, & Bloom, 2011; Hitchcock & Knobe, 2009; Knobe & Fraser, 2008; Kominsky et al., 2015). In situations in which two causal factors jointly cause an effect and one of them violates a prescriptive norm, participants tend to select this norm-violating factor over the other one as "the cause."

The *pen vignette* created by Knobe and Fraser (2008) illustrates this *norm effect*: "In a philosophy department, administrative assistants are allowed to take pens from the desk of the receptionist, whereas the faculty members are not allowed to do so. However, everyone takes pens regularly. When Professor Smith, a faculty member, and an administrative assistant simultaneously take pens one morning, a problem arises: there are no pens left." When the participants of the study were asked about who has "caused the problem" later, they tended to name Professor Smith rather than the administrative assistant although both agents took a pen, and therefore equally contributed to the problem. Since both agents merely differ in terms of their normative status (i.e., whether their behavior was right or wrong), examples like this are interpreted as showing that prescriptive norm violations can influence causal inferences.

These findings have motivated researchers to investigate the boundary conditions of this effect and to find an explanation. One important boundary condition of the norm effect seems to be that it is limited to causal selection judgments, that is, the selection of one primary cause in cases in which the combination of two causes is necessary for the target effect (*conjunctive causal structure*; see Kominsky et al., 2015). Interestingly, prescriptive norms do not influence intuitions about the structure or strength of the underlying causal model (Danks, Rose, & Machery, 2014; Samland & Waldmann, 2014). Other details and the explanation of the norm effect are still disputed. Whereas Hitchcock and Knobe (2009) assume that norm-violating causal factors are selected over norm-conforming ones because they provoke counterfactual thinking and serve as a better target of intervention, Alicke et al. (2011) ascribe the norm effect to a desire to blame a norm-violating agent which leads to an exaggeration of the agent's causal contribution to the outcome. Apart from these differences, most accounts agree that it is the *causal* judgment that is influenced by normative evaluations. There is, however, reason to doubt that the judgment in question is always a genuinely causal one.

The Accountability Hypothesis

What has been neglected in most discussions of the norm effect is the ambiguity of the term "cause." Depending on the context, the word "cause" can both refer to the question whether a mechanism underlying a causal relation is present and to the question whether an agent can be held *accountable* for an outcome. As Deigh (2008) points out, already Hart and Honoré (1959) have argued "(…) that the statement that someone has caused harm either means no more than that the harm would not have happened without ('but for') his action or (…) it is a disguised way of asserting the 'normative judgment' that he is responsible in the first sense, i.e., that it is proper or just to blame or punish him or make him pay" (pp. 61).

Both meanings of the term are inherent in the scenarios supporting the norm effect and can be represented as two hierarchically ordered layers of description of the presented causal relationship. The *causality layer* lies at the bottom and refers to the causal mechanisms connecting actions and outcomes in the scenarios. For example, in the pen vignette both agents initiate behaviors that lead to the removal of pens. The causality layer is a subset of the morally charged second layer, the *accountability layer*. Accountability assessments presuppose causality; people are only held

accountable for outcomes they have actually caused. However, additional components are required to assess accountability. In moral blame judgments, for example, the outcome needs to be negative (e.g., shortage of pens), and the agent needs to know the moral rule that forbids the act. Moreover, accountability increases (at least in Western societies) with intent.

The hierarchical structuring of the two layers of description of an action is crucial to the assumption that subjects often choose an accountability interpretation of the cause question over a merely causal one. Typically the causal mechanism described in the scenarios is trivial (e.g., taking pens). Moreover, all the additional features required for an accountability interpretation are explicitly or implicitly mentioned. Under these circumstances subjects might consider the accountability interpretation to be the intended meaning of the test question. Thus, a possible alternative explanation of the findings suggesting that causal inferences are influenced by norms might just be that subjects interpreted the test questions as a request to assess moral accountability.

Highlighting the Causal Meaning of Causal Test Questions

The accountability hypothesis assumes that subjects know that both agents cause the outcome. In accordance with this assumption, Samland and Waldmann (2014, Experiment 1) have shown that subjects do not differentiate between a norm-violating and a norm-conforming agent when a counterfactual test question is used that measures intuitions about *causal strength*. Both agents were considered to be equally causal according to this measure. However, norms could also affect *causal selection*. People may tend to choose one among many causes of an outcome based on pragmatic considerations, even if they are aware that all causes are equally important components of the mechanism. For example, the lighting of a match is typically picked as a cause of a fire in a forest, although most people know that oxygen also needs to be present (see Cheng & Novick, 1991). Accordingly, Hitchcock and Knobe (2009) have proposed that moral abnormality may highlight the counterfactual that the norm-violating agent might as well have behaved normally, which in turn triggers the intuition that this agent is the main cause. No such counterfactual is triggered for the norm-conforming agent, according to this theory.

While Hitchcock and Knobe's (2009) theory predicts that norm-violating agents who caused the outcome should be generally picked as the cause when contrasted with a norm-conforming co-agent, the accountability hypothesis predicts that pragmatic factors highlighting the accountability meaning of the test question are the main culprit. What are possible candidates for such pragmatic factors?

Studies supporting the influence of norms on causal selection typically share a specific structure: (i) the relevant causal information is *described* in a summarized presentation, and (ii) the potential causes are presented in a *personalized* way, that is, the question asks whether a specific person (e.g., Professor Smith) is the cause. Interestingly, each of these two features places emphasis on the accountability rather than only on the causality layer.

First, the transparent description of the causal setup is likely to create the impression that the causal component is a rather trivial part of the scenarios. That taking a pen removes a pen is of course known by all subjects so that they may conclude that the request to name the causes must refer to something different, namely accountability. Accordingly, Waldmann and Samland (2014, Experiment 2) could show that the norm-violating factor was no longer selected when the causal mechanism in the scenario was more complicated so that the causal component of the cover story became more salient than the accountability aspect. There was an alternative explanation of the findings, though. It could not be ruled out that subjects were more uncertain about the presented setup in the more complicated condition, and therefore did not differentiate between the two causes.

A better way to stress a causal understanding of the test question is therefore tested in the following Experiment 1: the presentation of the causal relationship in an experience-based manner. It is known that decisions that follow from descriptions can differ from decisions based on experience (e.g., Hertwig & Erev, 2009). In the judgment and decision-making literature, description-based decisions are based on summarized presentations of probabilities and outcomes (e.g., a 30% chance to win 5 $). By contrast, experience-based decisions are triggered by a series of individual learning trials in which information about individual instances (e.g., individual bets) is conveyed. Danks et al. (2014) have applied this distinction to norm-violation scenarios and have stated that "(…) people who learn from description typically engage in explicit, high-level reasoning that is slow, error-prone, and subject to outside influences. In contrast, those who learn from experience use other reasoning processes (…)" (p. 258). Consequently, outside influences like normative evaluations might be more likely to find their way into judgments based on descriptions. By contrast, trial-by-trial learning tasks in which subjects learn about a causal relation emphasize the relevance of the presented contingency. This should increase the tendency to interpret the test question as a causal one.

Second, most studies investigating the interaction of norms and causal judgments ask subjects to assess the causal relationship between a named agent and the outcome (e.g. "Professor Smith caused the problem."), although there are many elements in the causal chain between the agent and the outcome that constitute the causal process or mechanism. Each element could be the focus of a causal test question. For the accountability interpretation, however, the person is indeed the relevant factor since it is people who are held accountable for outcomes. Thus, it seems plausible that asking whether a person is a cause may suggest an accountability interpretation of the test question. This hypothesis is tested in Experiment 2 which demonstrates an

alternative way of how the interpretation of the test question can be shifted.

Experiment 1

Most experiments supporting an influence of norms on causal judgments have used the description format when presenting the scenarios. Since causal mechanisms are often trivial components of these scenarios, we hypothesized that subjects tend to interpret the test question as an attempt to assess accountability. One way to emphasize the causal mechanism is to use a trial-by-trial experience-based format to present the causal relation.

Danks et al. (2014) have tested the influence of moral norms on causal judgments in experience-based scenarios - but only for a causal relation with one single potential cause. The sequence of presented cases either exhibited a generative, a preventive or no causal relation and the potential cause was either a morally reprehensible or a neutral activity. Danks and colleagues could show that learners proved sensitive to the direction of the contingency but the morality of the causal factors did not influence the ratings. Since they did not present scenarios with two causal factors that are jointly responsible for the outcome (which is a prerequisite for the norm effect; see Kominsky et al., 2015), these results are only of limited use for our central focus on causal selection. A further shortcoming of the design of Danks et al. is that a control condition is missing in the reported experiment that demonstrates a norm effect in a described version of the chosen cover story.

In Experiment 1 we therefore contrasted a condition in which the causal relationship was described with a condition in which it was learned in a trial-by-trial learning phase. In both conditions, the same cover story was used in which two agents conjunctively caused a negative outcome. Our key question was whether a norm effect can be found irrespective of the condition, or whether it is restricted to the described scenario.

Theories that predict that the abnormality of the cause leads to a preferential selection (e.g., counterfactual theory; blame theory) should expect that the norm-violating cause should be selected regardless of the learning condition. Whether the causal structure is described or experienced should not make any difference. By contrast, the accountability hypothesis predicts that conveying the causal structure through a trial-by-trial contingency learning procedure would highlight the causality layer, whereas in the description condition the causality layer would be backgrounded. Thus, a norm effect is expected in the description but not in the contingency learning condition.

Method

Participants 86 undergraduates took part in the computer-based experiment that was run in a computer lab of the University of Göttingen. 14 subjects were excluded from the analysis because they failed to correctly answer a control question that checked whether the normative status of the two causes was understood properly. Thus, 72 subjects (84%) were included in the reported analyses (37 in the learning condition). The experiment was part of a battery of experiments; subjects earned 5 € for their participation.

Design The design of the experiment was based on a 2 (setting: learning vs. description) × 2 (normality: norm-violating vs. norm-conforming) structure with the last factor being manipulated within subject. Participants were randomly assigned to one of two conditions: a trial-by-trial contingency *learning* and a *description* condition.

Participants in both conditions were presented with a story about Tom who employs two gardeners, Alex and Benni. To foster growth, the gardeners have two chemicals at their disposal that can be used to protect plants against slugs and worms, A X200® and BOTANIX®. Since Tom has read that using the two different chemicals simultaneously can cause damage to the plants, he forbids the use of one of them (BOTANIX®). However, Benni, one of the gardeners, continues to use BOTANIX®. One day, Tom realizes that some of his plants are dried up which makes him feel miserable.

In the *description condition*, participants were then told that the plants that were harmed had grown in flower beds in which both gardeners had spread their chemicals: Alex had used the allowed chemical, whereas Benni had used the forbidden one.

In the *learning condition*, participants were given the same initial instruction as in the description condition but then were told that Tom would like to investigate the relationship between the use of the chemicals and the shriveled plants by conducting an empirical study. Subjects were then presented with a trial-by-trial learning phase in which they learned about the causal relationship between chemicals and plant growth. Subjects observed 10 slides in randomized order in which both chemicals were used and the plants were dried up, 20 slides in which only one of the two chemicals was applied (10 trials each) and the plants grew healthily, and 10 slides in which no chemical was sprayed so that no plant grew. Thus, the learning trials conveyed a conjunctive causal structure in which for healthy growth it was necessary that one of the causes was present, but not both.

Subsequently, participants in both conditions answered two causal questions about the chemicals ("How strongly did A X200®/BOTANIX® cause the plants' drying up?"). They expressed their judgment using an 11-point Likert-scale ranging from 0 (*not at all*) to 100 (*completely*). We used this test question and the rating scale because it is applicable to both learning conditions and our previous research has shown that subjects tended to give stronger ratings for the abnormal than the normal cause in a description condition (Samland & Waldmann, 2014).

Next, participants were asked the control question whether Alex and Benni had been allowed to use their chemical. Subjects who did not give the right answer were excluded from further analyses. On a last slide, we checked whether the conjunctive causal structure of the scenario was

understood. On a scale ranging from 0 to 100, participants were asked to specify the percentage of flower beds in which the plants were dried up when both, none, or one of the two chemicals were utilized. In the description condition, these four questions were introduced as hypotheticals because subjects in this condition did not see any data.

Results

Initially we checked whether subjects correctly understood the conjunctive causal model. Generally the level of understanding was very good. All participants in the learning condition correctly stated that shriveling occurred in 100 percent of the flower beds in which both chemicals had been applied and in 0 percent of those in which only one chemical had been used. Also the vast majority of subjects understood that the growth is absent when no chemical had been applied: the mean estimated percentage is close to 0 percent (M = 5.41, SD = 22.92). The ratings of participants in the description condition likewise indicated an understanding of the conjunctive causal model: the percentage of flower beds in which the plants dried up was estimated to be significantly higher if both chemicals had been used (M = 82.0, SD = 24.59) compared to cases in which only A X200® (M = 16.29, SD = 25.33), only BOTANIX® (M = 17.43, SD = 27.15) or no chemical (M = 19.14, SD = 26.61) had been applied.

The most important results concern the responses to the causal test questions (see Fig. 1). In the description condition, the forbidden chemical was given higher causal ratings compared to the chemical whose use had been allowed by Tom, $t(34) = 2.66$, $p = .01$. However, no significant difference between the ratings of the two chemicals was obtained in the learning condition in which the causal relationship between the two causes and the effect was presented in a trial-by-trial learning phase, $t(36) = 1.0$, $p = .32$.

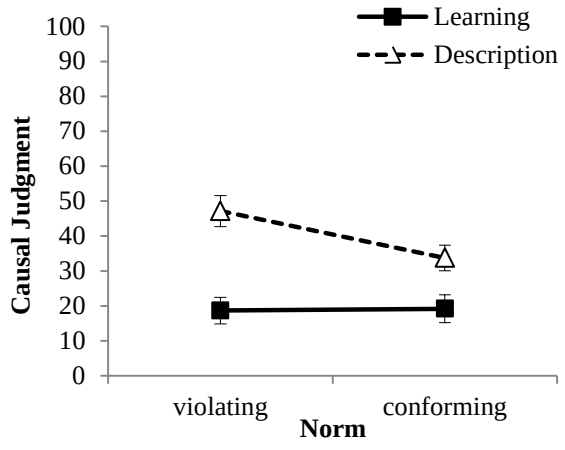

Figure 1: Results of Experiment 1. Error bars represent standard errors of means (*SE*).

An ANOVA yielded a main effect for normality, $F(1, 70) = 6.79$, $p = .01$, $\eta_p^2 = 0.09$, a main effect for setting, $F(1, 70) = 18.02$, $p < .001$, $\eta_p^2 = 0.20$, and, as predicted, a significant interaction between normality and setting, $F(1, 70) = 7.98$, $p = .006$, $\eta_p^2 = 0.10$.

Discussion

The experiment shows that the mode of presentation of causal information moderated the norm effect. Only in the description condition, which used the standard way of presenting the scenario, an effect of abnormality was observed. In this condition, the norm-violating cause was seen as more causal than the norm-conforming cause. However, this effect disappeared when, along with an initial description, learning trials were presented. This pattern is inconsistent with all theories that claim that abnormality determines causal selection regardless of the way causal information is presented. It is consistent with a pragmatic account that attributes the norm effect to a differential understanding of the test question. According to the accountability hypothesis, the low salience of the causal information and the emphasis on norms in the description condition may lead subjects to interpret the test question as a request to assess moral accountability of the agent. In the learning condition, on the other hand, the learning phase highlights the causal component of the scenario which may have led to the finding that the majority of subjects chose a causal interpretation of the test question.

One notable observation is that the norm effect in the description condition was relatively weak compared to previous studies. Samland and Waldmann (2014), for instance, presented their participants with three popular scenarios and used a similar test question with a scale ranging from 0 to 100 but the causal ratings for the norm-violating agent were, on average, higher than those in the present experiment. In previous experiments, the test question referred to the agents, for example, Professor Smith in the pen vignette. By contrast, in Experiment 1 we asked about chemicals. Although it is clear that chemicals by themselves do not generate harm but have to be applied by an agent, the role of the agent is still backgrounded. Given that people but not objects are typically held accountable, asking about the chemicals may have placed more emphasis on the causality layer and therefore contributed to the smaller effect. This hunch was explicitly tested in Experiment 2.

Experiment 2

Experiment 1 indicated that the understanding of the test question can be altered depending on whether the causal information was presented in a description or an experience format. Based on the findings of this experiment, we hypothesized that an additional factor that might influence the interpretation of the test question is the type of cause to which the question refers. Asking whether one of two persons is the cause should highlight the accountability interpretation because it is people not objects that are

typically blamed for aversive outcomes. By contrast, asking about the objects used in the actions might direct attention to the causal mechanism, and therefore emphasize the causality layer.

The scenario used in Experiment 2 is an adapted version of the pen vignette which we changed in three ways: First, we added an element in the causal chain from agent to outcome so that the agents could clearly be distinguished from the causal processes they initiate. In the present version the agents needed to press a colored button to get the requested office utensil. Since each button was only used by a single agent, the causal relation is the same, independent of whether we ask the causal test question about the person or the button. Nevertheless, we suspected that asking about a person would lead to a preference for an accountability interpretation, whereas asking about the button should emphasize the causal mechanism. Second, a third agent was introduced who did not contribute to the outcome and did not violate a norm. We introduced this agent to offer a candidate who is clearly non-causal. This should reduce the demand characteristic to differentiate between two equally causal agents, which might also contribute to a shift towards an accountability interpretation in the pen vignette. Third, we used a causal test question with categorical answer options instead of a rating scale to better capture the idea of causal selection. We allowed subjects to choose several agents as causal to avoid the demand characteristic that only one cause should be picked.

Method

Participants 213 subjects participated in the experiment that was run online in the U.K. 19 more people read the initial instructions but were not allowed to proceed in the experiment because they had failed a simple attention check. We excluded 74 (34.7%) participants who did not correctly remember which agent or button was norm-conforming or norm-violating. From the remaining set of 139 participants we excluded 41 (29.5%) additional subjects who did not correctly remember the causal relations described in the story.[1] We thus analyzed the data of 98 participants (57 in the person condition). Subjects were reimbursed with 50 British pence.

Design The design of the experiment was a 2 (question-type: person vs. mechanism) × 2 (normality: norm-violating vs. norm-conforming) structure with the last factor being manipulated within subject.

All participants were presented with a story about two departments in a philosophical faculty in which a chute system had been implemented so that office supplies can be delivered automatically into the offices of the employees. On each writing desk there are three differently colored buttons with which office supplies, such as pens or rubbers, can be ordered. The buttons are, however, not assigned to a specific product: each employee can individually program the assignment between color and office product. Participants then read that, due to a budget decision, only employees from department B are allowed to use the chute system to order office supplies. In contrast, employees from department A are no longer allowed to press the buttons located on their desks. Despite these new regulations, however, both departments continue ordering office supplies, which is frequently criticized by the receptionist. After these instructions, two attention questions were given asking how many buttons there are on each desk, and whether it is standardized which button leads to which office supply. Subjects who gave correct answers were presented with the following scenario:

> "One morning, by chance, Mrs. Smith from Department A, Mrs. Cooper from Department B and Mr. Wall from Department B press a button in their offices at the *exact same* time [09:26 a.m.]:
> - Mrs. Smith presses her **green** button and **a pen** is delivered to her office.
> - Mrs. Cooper presses her **blue** button and **a pen** is delivered to her office.
> - Mr. Wall presses his **yellow** button and **a rubber** is delivered to his office.
>
> A few minutes later [09:31 a.m.], the receptionist needs a pen and presses her pen button… but there are **no pens left** in the office supply store."

On the bottom of the page, the causal test question was shown: "What is the cause of the absence of pens in the office supply store?"

Participants were then randomly assigned to one of two conditions that differed only in the answer options describing the potential causes. In the *person condition*, participants answered the test question by ticking one or more of the following options: Mrs. Smith, Mrs. Cooper and Mr. Wall. The third agent, Mr. Wall, served as a control since he did not order any pens and he worked in the department that was allowed to order office supplies. In the *mechanism condition*, participants were presented with the test question and were offered the three options: green button, blue button and yellow button.

The causal judgment was followed by six comprehension and memory questions in which subjects were asked for each employee whether he or she was permitted to order office supplies by pressing a button (norm information) and what type of office supply he or she had ordered (causal structure information). Subjects who did not correctly answer the six comprehension questions were excluded from further analyses.

Results and Discussion

The results of the experiment are shown in Fig. 2. Overall, the distribution in the person condition is significantly different from the one in the mechanism condition, χ^2 (5, N = 98) = 14.52, p = .01. In the person condition, the norm-

[1] In our experience, drop out and exclusion rates depend on the online site. Typically, these numbers are lower in experiments run in the M-Turk community, which is not accessible to researchers outside the United States.

violating cause 1 (Mrs. Smith) was selected significantly more often as the single cause than the corresponding norm-violating cause 1 (green button) in the mechanism condition, χ^2 (1, N = 98) = 11.02, p < .001. Only in the person condition, the norm-violating cause 1 was selected significantly more frequently than the norm-conforming cause 2, χ^2 (1, N = 41) = 8.81, p = .003. No such effect was found in the mechanism condition, χ^2 (1, N = 17) = 0.06, p = .81. Furthermore, both causes together were selected more often in the mechanism condition than in the person condition, χ^2 (1, N = 98) = 4.91, p = .03.

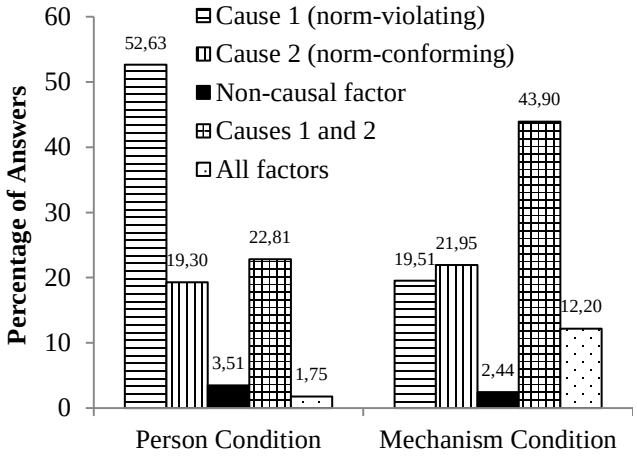

Figure 2: Results of Experiment 2. In the person condition, the two causes cause 1 and cause 2 refer to the story's agents Mrs. Smith and Mrs. Cooper, the non-causal factor to Mr. Wall. In the mechanism condition, cause 1 and cause 2 refer to the green and the blue button, whereas the yellow button is non-causal.

Thus, in the condition in which subjects chose among people as causes, a clear norm effect was replicated – the norm-violating cause 1 was selected over the norm-conforming cause 2. By contrast, no norm effect was found in the mechanism condition. Here the button corresponding to norm violation was selected equally often as the button corresponding to norm-conforming behavior and the dominant response was that both buttons are causal. These results support the hypothesis that in experiments in which subjects assess the causal status of people, the causal question is understood as a request to assess accountability. By contrast, buttons are not directly viewed as morally accountable and only indirectly refer to actions, which seems to foster a causal mechanism interpretation.

General Discussion

The present findings add to the evidence presented by Samland and Waldmann (2014) and further support the accountability hypothesis over alternative views (e.g., Alicke et al., 2011; Hitchcock & Knobe, 2009; Knobe & Fraser, 2008). Given the ambiguity of causal queries, we hypothesized that the influence of norms on causal judgments depends on aspects of the task that highlight an accountability interpretation rather than a causal understanding of the test question. This hypothesis was confirmed in two studies in which we manipulated the presentation mode of the causal information (description vs. contingency learning; Experiment 1) or varied whether the causal test question referred to a person or an object (Experiment 2).

Although our results indicate that norm effects in scenarios like the pen vignette are based on an accountability interpretation of the test question, it may still be that other factors, such as abnormality or covariation within the focal set, are motivating causal selection in other contexts (Cheng & Novick, 1991; Hitchcock & Knobe, 2009; Kominsky et al., 2015). Future research will have to test the boundary conditions of different possible mechanisms of causal selection.

References

Alicke, M. D. (1992). Culpable causation. *Journal of Personality and Social Psychology*, *63*, 368-378.

Alicke, M. D., Rose, D., & Bloom, D. (2011). Causation, norm violation, and culpable control. *Journal of Philosophy*, *108*, 670-696.

Cheng, P. W., & Novick, L. R. (1991). Causes versus enabling conditions. *Cognition*, *40*, 83-120.

Danks, D., Rose, D., & Machery, E. (2014). Demoralizing causation. *Philosophical Studies*, *171*, 251-277.

Deigh, J. (2008). Can you be morally responsible for someone's death if nothing you did caused it? In W. Sinnott-Armstrong (Hrsg.), *Moral Psychology* (pp. 449-461). Massachusetts: MIT Press.

Hart, H. L. A., & Honoré, A. M. (1959). *Causation in the law*. Oxford University Press, Oxford.

Hertwig, R., & Erev, I. (2009). The description-experience gap in risky choice. *Trends in Cognitive Sciences*, *13*, 517-523.

Hitchcock, C., & Knobe, J. (2009). Cause and norm. *Journal of Philosophy*, *106*, 587-612.

Kominsky, J. F., Phillips, J., Gerstenberg, T., Lagnado, D. A., & Knobe, J. (2015). Causal superseding. *Cognition*, *137*, 196-209.

Knobe, J., & Fraser, B. (2008). Causal judgment and moral judgment: Two experiments. In W. Sinnott-Armstrong (Ed.), *Moral Psychology* (pp. 441-447). Massachusetts: MIT Press.

Samland, J., & Waldmann, M. R. (2014). Do social norms influence causal inferences? In P. Bello, M. Guarini, M. McShane & B. Scassellati (Eds.), *Proceedings of the 36th Annual Conference of the Cognitive Science Society* (pp. 1359-1364). Austin, TX: Cognitive Science Society.