

BRIEF REPORT

Expectations Affect Physical Causation Judgments

Tobias Gerstenberg and Thomas Icard
Stanford University

When several causes contributed to an outcome, people often single out one as “the” cause. What explains this selection? Previous work has argued that people select abnormal events as causes, though recent work has shown that sometimes normal events are preferred over abnormal ones. Existing studies have relied on vignettes that commonly feature agents committing immoral acts. An important challenge to the thesis that norms permeate causal reasoning is that people’s responses may merely reflect pragmatic or social reasoning rather than arising from causal cognition per se. We tested this hypothesis by asking whether the previously observed patterns of causal selection emerge in tasks that recruit participants’ causal reasoning about physical systems. Strikingly, we found that the same patterns observed in vignette studies with intentional agents arise in visual animations of physical interactions. Our results demonstrate how deeply normative expectations affect causal cognition.

Keywords: causality, counterfactuals, causal selection, normality, intuitive physics


Supplemental materials: <http://dx.doi.org/10.1037/xge0000670.supp>

In 2015, the U.S. Supreme Court affirmed the constitutionality of same-sex marriage across the United States with a vote of five to four (Kennedy, 2015). Each of the five judges who voted in favor was clearly “a cause” of the outcome; after all, the outcome depended on each of them voting the way they did. However, the vote of one justice struck many as “the cause” of the outcome: Among the five justices who were in favor of same-sex marriage, only one, Justice Anthony Kennedy, was appointed by a Republican president.

How do people select “the” cause of an outcome when multiple causes contributed? The problem of causal selection has long puzzled philosophers; it is fundamental in the law, for example, for the distinction between factual and legal cause (Hart & Honoré, 1985; Stapleton, 2001, 2008; Summers, 2018), and it has received renewed attention in the psychological sciences in recent years (Hitchcock & Knobe, 2009; Icard, Kominsky, & Knobe, 2017; Kirfel & Lagnado, 2018; Kominsky, Phillips, Gerstenberg, Lagnado, & Knobe, 2015; Samland & Waldmann, 2016).

As the example just described illustrates, one key factor that influences causal selection is the perceived normality of each cause. People tend to select abnormal (or unexpected) over normal (or expected) causes (Cheng & Novick, 1991; Hesslow, 1988; Hilton & Slugoski, 1986; Knobe & Fraser, 2008). Norms and expectations also play a critical role in how people make causal judgments about omissions, that is, events that didn’t happen. At any given moment in time, an infinite number of events do not happen, but people normally only ever consider those that violated their expectations (Clarke, Shepherd, Stigall, Waller, & Zarpentine, 2015; Gerstenberg & Tenenbaum, 2016; Henne, Niemi, Pinillos, De Brigard, & Knobe, 2019; McGrath, 2005; Samland & Waldmann, 2016; Stephan, Willemsen, & Gerstenberg, 2017; Willemsen, 2016; Wolff, Barbey, & Hausknecht, 2010, 2011).

Different explanations for this abnormal selection effect have been proposed. Some have argued that abnormal events are preferentially selected because they spontaneously trigger counterfactual thoughts about the normal course of events (Kahneman & Miller, 1986; Kahneman & Tversky, 1982; Petrocelli, Percy, Sherman, & Tormala, 2011; Phillips, Luguri, & Knobe, 2015; Schaffer, 2005). Others have suggested that the focus on abnormal events results from pragmatic principles of communication (Grice, 1975; Hilton, 1990). When telling someone about what happened and why, it is more important to mention unexpected events because people can infer the expected events themselves. Further proposals have argued that causal selections are influenced by moral concerns such as who is to blame (Alicke, 2000; Samland & Waldmann, 2016) or to be held responsible (Sytsma, Livengood, & Rose, 2012) and that the agents’ epistemic states matter (Kirfel & Lagnado, 2019).

 Tobias Gerstenberg, Department of Psychology, Stanford University; Thomas Icard, Department of Philosophy, Stanford University.

Part of this research was presented at the 2017 Annual Meeting of the Society of Philosophy and Psychology at Johns Hopkins University, Baltimore, Maryland. Tobias Gerstenberg was supported by the Center for Brains, Minds & Machines, which is funded by the National Science Foundation’s Science and Technology Center (Award CCF-1231216). We thank Jonathan Livengood, Jonathan Phillips, Kevin Smith, and Joshua Tenenbaum for valuable feedback.

Correspondence concerning this article should be addressed to Tobias Gerstenberg, Department of Psychology, Stanford University, 302 Jordan Hall, Stanford, CA 94305. E-mail: gerstenberg@stanford.edu

Although many existing efforts have been aimed at explaining why abnormal causes tend to be selected, recent work has shown that people sometimes select a normal rather than an abnormal event as the cause of the outcome (Gavanski & Wells, 1989; Gerstenberg et al., 2018; Icard et al., 2017; Johnson & Rips, 2015; Sytsma et al., 2012). For example, Icard et al. (2017) presented participants with a story in which Suzy and Billy both enter a room and set off a motion detector. The experiment manipulated whether the two causes were both necessary (the motion detector is triggered only if more than one person enters the room) or individually sufficient (the motion detector is triggered whenever someone enters the room) for bringing about an outcome. Further, the authors manipulated whether Billy was told to arrive at the same time as Suzy or told not to come at that time. The results showed that when Billy's action was unexpected and both causes were required, participants tended to agree that Billy caused the outcome, whereas when only one of the causes was needed, participants tended to disagree that Billy caused the outcome.

To date, there is no unified account that accurately captures all of the known patterns of normality effects on causal selection. Some accounts have predicted that abnormal events are generally favored (Hall, 2007; Halpern, 2016; Halpern & Hitchcock, 2015). Others have predicted that normal events are held more responsible for positive outcomes (Johnson & Rips, 2015), that people select normal causes for normal outcomes and abnormal causes for abnormal outcomes (Harinen, 2017; see also Gavanski & Wells, 1989), that it depends on the causal structure (Icard et al., 2017; Morris, Phillips, Gerstenberg, & Cushman, 2019), or that it depends on what one infers about the actors from their actions (Gerstenberg et al., 2018).

We do not develop a unified account of causal selection here. Instead, we provide critical data to inform the debate by showing that the phenomenon is much more pervasive than might have been expected. All empirical work on causal selection so far has relied on vignettes, which nearly always feature intentional agents.¹ The crucial information about norms is always communicated explicitly by telling participants that some agent was "not supposed to" perform some action or that an event was "unlikely" to occur. Although such manipulations have been shown to affect people's judgments in these cases, it is unclear to what extent these effects may be due to pragmatic and social reasoning (see, e.g., Alicke, 2000; Samland, Josephs, Waldmann, & Rakoczy, 2016; Samland & Waldmann, 2016; Sytsma et al., 2012) or whether they reveal something deeper about causal cognition more generally. For instance, Danks, Rose, and Machery (2014) questioned whether previous results would extend beyond vignette-based experiments:

The existing body of evidence may well show that certain kinds of causal cognition—specifically, a highly language-driven kind of causal reasoning—are influenced by normative considerations, but it falls short of showing that causal cognition in general is influenced by normative considerations. (p. 256)

Is the influence of norms limited to vignettes and linguistically described scenarios, especially those that involve agents performing immoral actions? Or should one expect to see this influence throughout causal cognition? To help answer this question, we studied causal judgments involving inanimate objects, based on physically realistic animations of collision events (see Gerstenberg, Goodman, Lagnado, & Tenenbaum, 2012, 2014, 2015; Ger-

stenberg, Peterson, Goodman, Lagnado, & Tenenbaum, 2017; Gerstenberg & Tenenbaum, 2017; Stephan et al., 2017; Wolff, 2007). In addition to being a very different setting from those employed in previous work on this phenomenon, the visual nature of the stimuli brings three further experimental advantages.

First, statistical norms can be directly experienced: Rather than being told about norms explicitly, participants in our experiments were presented with sequences of videos, allowing them to form probabilistic expectations about the various events directly from their observations. It has been shown that people make different decisions depending on whether information is directly experienced or merely described (see, e.g., Hertwig, Barron, Weber, & Erev, 2004). Our setup allowed us to confirm that the way in which statistical norms are presented does not change previously observed patterns.

Second, we enjoyed much greater control over participants' interpretations of the events, because they observed exactly what happens across the entire scenario. Our videos showed that each cause contributed identically to the outcome—they differed only in their prior probabilities.

Third, the format allowed flexible manipulation of both causal structure and probabilistic information. We took advantage of this by going beyond the causal structures that have been studied in previous work.

To the extent that norms affect participants' causal selections even in this visual setting with inanimate objects—and especially to the extent that people's judgments show the very same patterns as in existing vignette studies—this provides strong evidence that the influence of norms on causal selection is general and systematic and that it occurs even in the absence of explicit social or moral concerns.

Experiment

Past research has shown that people's causal selections in vignette studies are affected by both normative expectations and the causal structure of the situation (Icard et al., 2017; Kominsky et al., 2015). We studied how manipulating normality and causal structure between two candidate causes across a range of physically realistic animations affects participants' causal selections. Whereas prior research has looked at how different kinds of norms affect causal selections, such as prospective norms about how someone should act (Icard et al., 2017; Kominsky et al., 2015; Samland & Waldmann, 2016) or norms of proper functioning that dictate how a device should work (Hitchcock & Knobe, 2009), we focused here on the role of statistical norms (see also Kirfel & Lagnado, 2018).

Figure 1 shows some examples of the setup. The animations feature three balls, A, B, and E. Ball E is initially stationary, and Balls A and B enter the scene from the right. The outcome of interest is whether or not Ball E goes through the gate on the left. There are two blockers that sometimes block A's and B's motion and sometimes let them through. We manipulated the probability

¹ We are aware of only one vignette that didn't feature agents. Hitchcock and Knobe (2009) described a machine that is set up in a way such that it will short circuit if both a black wire and a red wire touch a battery at the same time. The black wire is supposed to touch the battery, whereas the red wire isn't. When both wires in fact touched the battery, participants agreed more with a statement that the red rather than black wire touching the battery caused the machine to short circuit.

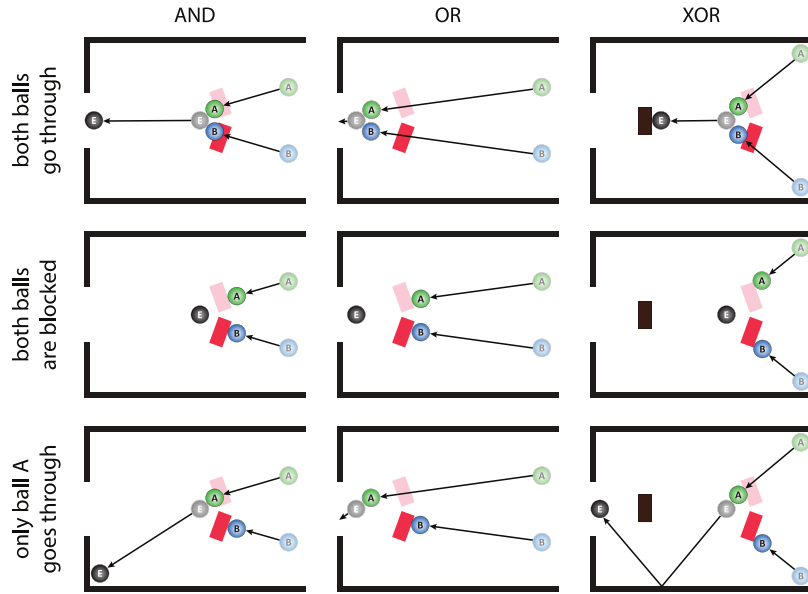


Figure 1. Diagrammatic illustration of clips in the AND, OR, and XOR condition of the experiment. The color (shade of gray in the print version of the article) of the blocker indicates its probability of blocking a ball. The dark red (dark gray) blocker has an 80% chance of blocking a ball, and the light red (light gray) blocker has a 20% chance of blocking a ball. The top row shows cases in which both balls went through the blocker. The middle row shows cases in which neither ball went through the blocker. The bottom row shows cases in which only Ball A went through the blocker (videos of the stimuli may be viewed at https://github.com/tobiasgerstenberg/causal_expectations). See the online article for the color version of this figure.

that a given blocker worked, thus making it more or less (statistically) normal for a ball to go through the blocker. In some situations, both Ball A and Ball B needed to go through the blocker to make Ball E go through the gate (AND structure), whereas in other situations, either ball was individually sufficient (OR structure), or Ball E only went through the gate if exactly one of the balls went through the blocker (XOR structure). Finally, in the test clip for which participants were asked to select the cause, we manipulated whether both Balls A and B went through the blocker (see Figure 1, top row) or both were blocked (see Figure 1, middle row).

Predictions

The primary goal of this study was to investigate how statistical expectations as well as the causal structure of the situation affect causal selections. By using visual animations of physical interactions rather than written vignettes, we minimized any uncertainty participants may have had about what actually happened. This way, we could make sure that any potential differences between participants’ causal selections were due to only the manipulation of statistical norms and the causal structure of the situation.

Though comparing theoretical proposals was not our main focus, it was useful to consider existing accounts aimed at accounting for previous studies. Here, we briefly describe what patterns of causal selections one should expect based on several accounts from the literature.²

No Effect of Normality

One possibility is that manipulating statistical normality has no effect on participants’ causal selections in visually presented stimuli (cf. Danks et al., 2014). Our videos clearly showed that each cause contributed to the outcome in exactly the same way. For example, when both balls went through the blockers (see Figure 1, top row), Balls A and B collided with Ball E in a symmetrical fashion and each ball had the same causal effect on Ball E’s subsequent movement. If participants focused on just each cause’s actual contribution to the outcome, their causal selections should be at chance.

Abnormal Selection

The abnormal selection hypothesis (Halpern & Hitchcock, 2015; Hart & Honoré, 1985; Hilton & Slugoski, 1986; Hitchcock & Knobe, 2009) predicts that participants will always select the event that was less expected to happen as the cause of the outcome.

Correspondence

According to the correspondence hypothesis (Gavanski & Wells, 1989; Harinen, 2017), participants select normal causes for normal effects and abnormal causes for abnormal effects. For example, in the AND structure, Ball E only rarely goes through the

² In the online supplemental materials, we describe in more detail how the quantitative predictions were derived.

gate. So, if Ball E does go through the gate (abnormal effect), participants are predicted to select the abnormal event as the cause (i.e., the ball that used to get blocked in the past). In contrast, in the OR structure, Ball E almost always goes through the gate. Here, Ball E's going through the gate is a normal effect, so participants are predicted to select the normal event as the cause (i.e., the ball that normally goes through the motion block).

Necessity and Sufficiency

Several researchers have argued that an event's causal status depends on the extent to which it was necessary and sufficient for bringing about the outcome (Fitelson & Hitchcock, 2011; Gerstenberg et al., 2015; Icard et al., 2017; Kominsky et al., 2015; Lagnado, Gerstenberg, & Zultan, 2013; McGill & Tenbrunsel, 2000; Pearl, 1999). Existing theories differ in how exactly *necessity* and *sufficiency* are defined, as well as in how both notions combine to affect causal judgments. For concreteness we followed Icard et al. (2017), who assumed that causal strength is a weighted function of "actual" necessity (i.e., the extent to which the cause was necessary in the given situation) and "robust" sufficiency (i.e., the extent to which the cause is generally sufficient for bringing about the outcome). The relative weighting is determined by the perceived normality of the cause. Because a cause's necessity and sufficiency are affected by the causal structure, this account predicts (like the correspondence hypothesis) that both normality and causal structure will affect participants' selections.

Method

Participants

In our study, 300 participants ($N_{\text{female}} = 137$; $M_{\text{age}} = 35$, $SD = 11$) were recruited via Amazon Mechanical Turk (Crump, McDonnell, & Gureckis, 2013) using psiTurk (Gureckis et al., 2016). The study was approved by the institutional review board. An additional 99 participants were excluded because they failed one or more of the exclusion criteria specified later. We imposed stringent exclusion criteria to make sure that only those participants who correctly encoded the causal structure as well as the statistical information were included. We collected participants until we had 50 participants in each of the six conditions who passed all of the exclusion criteria. The power analysis tool G3.1 (Faul, Erdfelder, Lang, & Buchner, 2007) showed that 49 participants would be required to detect an effect of $g = .25$ with power $1 - \beta = .95$ and $\alpha = .05$.

Design and Procedure

The experiment has a 3 (causal structure: AND, OR, XOR) \times 2 (test clip: both balls go through, both balls are blocked) design.³ Both factors were manipulated between participants (with $N = 50$ participants in each condition). The positions of the balls and blockers were counterbalanced between participants. This means that for some participants Ball A was more likely to go through the blockers and that for others Ball B was more likely.

Participants first received instructions about the task. They were told that they would see a number of video clips in which Balls A and B would always start on the right side and Ball E would be

initially at rest. They were informed about the blockers (as well as the solid block in the XOR condition) and that darker blockers would be more likely to block balls than would lighter ones. Three diagrams illustrated the general setup. One diagram showed the initial setting, another showed what would happen if only one ball went through the blocker, and the third one showed what would happen when both balls went through the blocker. In addition, participants were instructed about whether Ball E would go through the gate if either or both Ball A and B went through their respective blockers. Finally, participants learned that they would see 10 video clips and that they should pay close attention to what happens in each clip. In particular, they should try to get a sense for how good each blocker is.

Table 1 summarizes what clips participants saw and how often. After each clip, participants were asked to answer three yes–no forced-choice questions (shown simultaneously in this order): (a) "Did ball A get blocked?" (b) "Did ball B get blocked?" and (c) "Did ball E go through the gate?" Participants were allowed to proceed to the next video clip only after having answered each question. These questions ensured that participants were paying close attention to what happened in each clip.

After watching the 10 clips, participants were told that they would now see one more clip and that they would be asked different questions about this clip. They were told that they would first see only the beginning of the clip and be asked to make a prediction about what would happen. They would then see the clip until the end and answer a question about what did happen. In the test clip, Balls A and B started moving and then the clip paused. Participants were asked to what extent they agreed with the following three statements: (a) "Ball A will hit ball E," (b) "Ball B will hit ball E," and (c) "If only one of the two balls goes through the block and hits ball E then ball E will go through the gate." All three questions were shown simultaneously, and participants provided their responses on sliding scales where the end points were labeled 0 and 100.

Finally, participants saw the clip played until the end. Depending on the experimental condition, either both balls went through the blocker (see Figure 1, top row) or both balls got blocked (see Figure 1, middle row). After the end of the clip, participants were asked to select which of the following two statements better described what happened: "Ball E went through the gate because ball A [ball B] went through the motion block." The statements were adapted based on the experimental condition (e.g., another statement would be: "Ball E didn't go through the gate because ball A [ball B] didn't go through the motion block").

Participants completed the experiment by providing their demographic information and by answering a question about what factors influenced their causal judgment in a free-form text box. On average, the experiment took 6 min ($SD = 2.82$) to complete.

Exclusion Criteria

Based on a pilot experiment, we defined the following set of exclusion criteria. A participant was excluded if any of these were true:

³ All materials, including data, analysis scripts, experiments, figures, and videos, may be accessed at https://github.com/tobiasgerstenberg/causal_expectations.

Table 1
Contingency Table Illustrating What Happened in the First 10 Video Clips That Participants Saw

| | Ball B goes through | Ball B gets blocked | Ball A overall |
|---------------------|---------------------|---------------------|----------------|
| Ball A goes through | 1 | 7 | 8 |
| Ball A gets blocked | 1 | 1 | 2 |
| Ball B overall | 2 | 8 | 10 |

Note. This table illustrates the case in which Ball A has an 80% chance of going through the blocker and Ball B has a 20% of going through. In the experiment, we counterbalanced which ball was more likely to go through the blocker.

- The participant answered fewer than 28 out of 30 questions correctly about what happened in the 10 test clips.
- When asked to agree with the statements about whether Ball A would hit Ball E, and whether Ball B would hit Ball E, the judgment for the 80% ball was less than 20 points greater than the judgment for the 20% ball (the sliders ranged from 0 to 100).
- When asked to agree whether Ball E would go through the gate if only one of the two balls went through the blocker, the participant’s answer was on the incorrect side of the scale (i.e., less than 50 for OR and XOR, and greater than 50 for AND).

These strict exclusion criteria ensured that only participants who correctly encoded the statistical and causal relationships were included in the final data set.

Results

Figure 2 shows the results for all six conditions in the experiment together with the predictions of three different accounts we introduced earlier. Each bar shows the percentage

with which participants selected either of the two balls. The three bars on the left side show the results when both balls went through the gate in the test clip, and the bars on the right side show the results when both balls were blocked. The labels on the x-axis indicate the different causal structures. We report the percentage of participants who selected the ball that normally gets blocked, together with the confidence intervals of a two-sided binomial test that compares the observed proportion to an expected proportion of 50%, as well as Bonferroni-adjusted *p* values.

When both balls went through and the causal structure was AND (cf. Figure 1, top left), participants overwhelmingly selected the ball that normally gets blocked (96% confidence interval [CI: 86%, 100%], *p* < .001). However, when the structure was OR (cf. Figure 1, top middle), participants said that Ball E went through the gate because of the ball that normally goes through the blocker (22% CI [12%, 36%], *p* < .001). Even though each ball’s actual contribution to the outcome was identical in the two causal structures, participants’ tendency to select one cause over the other was reversed between the AND and OR structures. In the XOR structure (cf. Figure 1, top right), Ball E didn’t go through the gate when both Balls A and B went through the blocker. Here, participants selected that Ball E didn’t go through the gate because of the ball that normally gets blocked (94% CI [83%, 99%], *p* < .001).

When both balls were blocked, in the AND structure, participants were divided about why it was that Ball E didn’t go through the gate (54% CI [39%, 68%], *p* = 4.031). For both the OR and XOR structures, participants overwhelmingly said that Ball E didn’t go through the gate because of the ball that normally goes through the blocker but didn’t do so this time (12% CI [5%, 24%], *p* < .001, and 8% CI [2%, 19%], *p* < .001, respectively).

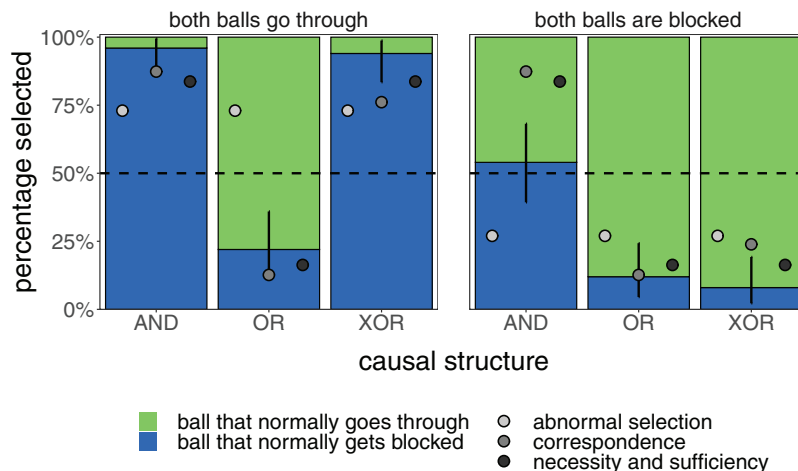


Figure 2. Percentage of participants who selected the ball that normally gets blocked (blue; darker gray in the print version of the article) or the ball that normally goes through (green; lighter gray) separated by the causal structure (x-axis), as well as whether both balls go through in the test clip (left panel) or both balls are blocked (right panel). The error bars indicate 95% confidence intervals. The dashed line indicates the chance level. The white, gray, and black dots indicate the different accounts’ predictions. See the online article for the color version of this figure.

Overall, when both cause events were necessary for the outcome event to occur, participants had a strong tendency to select the abnormal cause. This was the case for Ball E going through in the AND structure and for Ball E not going through in the OR structure (for which it was necessary for both Balls A and B to get blocked).

In situations in which the outcome was overdetermined, participants' selections were less clear. When Ball E went through the gate in the OR structure (where either ball would have been sufficient for making E go through the gate), participants favored the normal cause. However, when Ball E didn't go through the gate in the AND structure (where either ball's not going through the block was sufficient for the negative outcome), participants were equally likely to choose either of the two causes. In the XOR structure, participants strongly preferred the abnormal cause as the reason for why Ball E didn't go through the gate.

Overall, we found that despite the fact that both balls' causal contributions in the test clip were identical, participants displayed clear preferences for selecting one cause over another. Participants' selections are not well explained by an *abnormal selection account*, which predicts that the abnormal cause would be selected in each situation. Instead, participants' selections were affected by both normality and causal structure in a way that is largely consistent with the predictions of both the *correspondence account* and the *necessity and sufficiency account* (see Figure 2). Both of these accounts correctly predict that participants will be more likely to select the *normal* cause in situations in which the outcome was causally overdetermined, compared to situations in which both causes were necessary for the outcome to occur. These patterns, in turn, are consistent with those in previous empirical work based on vignettes (e.g., Icard et al., 2017; Kominsky et al., 2015). However, none of the theoretical accounts predicts the full pattern of results. Specifically, when both balls are blocked in the AND structure, participants displayed no preference in their causal selection, whereas both the correspondence account and the necessity and sufficiency account predict that the ball that normally gets blocked will be selected.

Discussion

How do people choose which of several causes deserves the status of "the" cause? A wealth of research has shown that judgments about this question are influenced by norms involving the various candidate causes. So far, this research has almost exclusively relied on vignette studies (Hilton & Slugoski, 1986; Icard et al., 2017; Knobe & Fraser, 2008; Kominsky et al., 2015; Samland & Waldmann, 2016; although see Kirfel & Lagnado, 2018, 2019). Here, we show for the first time that normality affects participants' causal selections even when scenarios are presented as physically realistic visual animations. These animations minimize any potential uncertainty about what actually happened and mitigate any role for the kinds of social or moral concerns that have been argued to drive causal selections.

In particular, the results replicate and extend the striking interaction between normality and causal structure that was first found in vignette studies (Icard et al., 2017; Kominsky et al., 2015). People do not always select the abnormal cause. When two causes bring about an outcome, people select the abnormal cause in an AND structure but select the normal cause in an OR

structure. Even though participants see that both causes affected the outcome in identical fashion, statistical normality has a strong and systematic effect on causal selection. The same pattern of causal selection was observed, although less pronounced, for judgments of omissive causation. Finally, the results show that in XOR structures, participants strongly preferred to select the abnormal cause.

This intriguing interaction of normality and causal structure is largely consistent with the predictions of at least two existing accounts. First, correspondence theory (Harinen, 2017) predicts that people select abnormal causes for abnormal effects and normal causes for normal effects. Second, Icard et al.'s (2017) account predicts that causal selections are sensitive to how necessary and sufficient each candidate cause was for bringing about the outcome.

Another style of account—potentially consistent with those already discussed but conceptually distinct—is that participants select causes that provide good points for intervention (cf. Hitchcock, 2012; Lombrozo, 2010; Morris et al., 2018; Woodward, 2006). Imagine that a person wants to make Ball E go through the gate and can intervene to increase the probability of either of the two events happening (Eaton & Murphy, 2007; Eberhardt & Scheines, 2007). In an AND structure, it makes sense to increase the probability of the less likely cause (i.e., the ball that normally gets blocked). In an OR structure, one should intervene on the more likely cause (i.e., the ball that normally goes through). Thus, this *optimal intervention account* also predicts the observed reversal effect. Further experiments are needed to tease these different accounts apart, for example by manipulating the normality of the two candidate causes across a wide range of settings.

Although the correspondence account and the necessity and sufficiency account capture the overall pattern of results well, there are two trends in the data that aren't captured by either account. First, participants' tendency to select one cause over the other was less pronounced when the outcome was overdetermined (i.e., in the OR structure when both balls went through, and in the AND structure when both balls were blocked). Note that it is in these situations of overdetermination that the three accounts make qualitatively different predictions. It is plausible that some participants generally prefer to select abnormal causes irrespective of the causal structure, whereas other participants consider both normality and causal structure as dictated by the other accounts, leading to the observed aggregation effect.

Second, participants' intuitions about which cause to select were split when the outcome was overdetermined and both of the balls were blocked—and more so than in the logically equivalent situation in which both balls went through. This asymmetry suggests that the effect of causal structure and normality may play out differently when the event of interest is an omission (Spranca, Minsk, & Baron, 1991).

Our results show that manipulating statistical norms and causal structure is sufficient to produce a reversal effect in people's causal selections. We haven't ruled out that more general normative considerations may have affected the results. For example, participants may have attributed agency to the objects in our simulations (Csibra, 2008; Scholl & Tremoulet, 2000), construed the outcome of Ball E going through the gate

as positive, and considered who was deserving of praise for the positive outcome. In future work, we plan to manipulate outcome valence to see whether it affects participants' causal selections.⁴

Our results also suggest that the role of norms in causal cognition is more pervasive than one might have previously thought. Normative expectations affect participants' causal judgments not only about agents in written vignettes but also in visual displays of physical interactions. These results are relevant for the theoretical debate surrounding causal judgment more generally. The systematic effect of norms on causal selections presents a challenge for theories predicting that causal judgments are grounded solely in the actual processes that unfolded (e.g., Wolff, 2007). In contrast, theories that link causal judgment and causal selection to the consideration of counterfactuals can more readily incorporate the role of norms: Normative considerations affect what counterfactuals come to mind, which in turn determine what causal judgments people reach (Gerstenberg et al., 2017; Hitchcock & Knobe, 2009; Kahneman & Miller, 1986; McGrath, 2005).

Whereas we have focused here specifically on the role of statistical norms for causal selection, the question of how deep norms affect causal cognition more generally remains an important area for future research (Danks et al., 2014; Phillips et al., 2015).

Conclusion

When multiple causes bring about an outcome, people's judgments about which event was "the" cause of the outcome depend systematically on causal structure. In some cases abnormal events are selected, in other cases normal events are selected. Whereas previous studies have used vignettes to show that causal structure and normality interact to produce different causal judgments (Icard et al., 2017; Kominsky et al., 2015), the experiment reported here is the first to demonstrate this intriguing effect using visual animations. Even when one can directly see that both causes contributed to the outcome in a perfectly symmetric fashion, and when there are no moral or social reasons for favoring one cause over the other, normative expectations still affect which one a person selects as the cause. Norms influence causal cognition in a deep and systematic way, revealing distinctive patterns that have yet to receive a comprehensive explanation.

⁴ We thank Jonathan Livengood for suggesting this manipulation.

References

- Alicke, M. D. (2000). Culpable control and the psychology of blame. *Psychological Bulletin*, *126*, 556–574. <http://dx.doi.org/10.1037/0033-2909.126.4.556>
- Cheng, P. W., & Novick, L. R. (1991). Causes versus enabling conditions. *Cognition*, *40*, 83–120. [http://dx.doi.org/10.1016/0010-0277\(91\)90047-8](http://dx.doi.org/10.1016/0010-0277(91)90047-8)
- Clarke, R., Shepherd, J., Stigall, J., Waller, R. R., & Zarpentine, C. (2015). Causation, norms, and omissions: A study of causal judgments. *Philosophical Psychology*, *28*, 279–293. <http://dx.doi.org/10.1080/09515089.2013.815099>
- Crump, M. J. C., McDonnell, J. V., & Gureckis, T. M. (2013). Evaluating Amazon's Mechanical Turk as a tool for experimental behavioral research. *PLoS ONE*, *8*(3), e57410. <http://dx.doi.org/10.1371/journal.pone.0057410>
- Csibra, G. (2008). Goal attribution to inanimate agents by 6.5-month-old infants. *Cognition*, *107*, 705–717. <http://dx.doi.org/10.1016/j.cognition.2007.08.001>
- Danks, D., Rose, D., & Machery, E. (2014). Demoralizing causation. *Philosophical Studies*, *171*, 251–277. <http://dx.doi.org/10.1007/s11098-013-0266-8>
- Eaton, D., & Murphy, K. (2007). Exact Bayesian structure learning from uncertain interventions. In M. Meila & X. Shen (Eds.), *Proceedings of machine learning research: Vol. 2. Artificial intelligence and statistics* (pp. 107–114). Retrieved from <http://proceedings.mlr.press/v2/eaton07a/eaton07a.pdf>
- Eberhardt, F., & Scheines, R. (2007). Interventions and causal inference. *Philosophy of Science*, *74*, 981–995. <http://dx.doi.org/10.1086/525638>
- Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2007). G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, *39*, 175–191. <http://dx.doi.org/10.3758/BF03193146>
- Fitelson, B., & Hitchcock, C. (2011). Probabilistic measures of causal strength. In P. McKay Illari, F. Russo, & J. Williamson (Eds.), *Causality in the sciences* (pp. 600–627). <http://dx.doi.org/10.1093/acprof:oso/9780199574131.003.0029>
- Gavanski, I., & Wells, G. L. (1989). Counterfactual processing of normal and exceptional events. *Journal of Experimental Social Psychology*, *25*, 314–325. [http://dx.doi.org/10.1016/0022-1031\(89\)90025-5](http://dx.doi.org/10.1016/0022-1031(89)90025-5)
- Gerstenberg, T., Goodman, N. D., Lagnado, D. A., & Tenenbaum, J. B. (2012). Noisy Newtons: Unifying process and dependency accounts of causal attribution. In N. Miyake, D. Peebles, & R. P. Cooper (Eds.), *Proceedings of the 34th Annual Conference of the Cognitive Science Society* (pp. 378–383). Austin, TX: Cognitive Science Society.
- Gerstenberg, T., Goodman, N. D., Lagnado, D. A., & Tenenbaum, J. B. (2014). From counterfactual simulation to causal judgment. In P. Bello, M. Guarini, M. McShane, & B. Scassellati (Eds.), *Proceedings of the 36th Annual Conference of the Cognitive Science Society* (pp. 523–528). Austin, TX: Cognitive Science Society.
- Gerstenberg, T., Goodman, N. D., Lagnado, D. A., & Tenenbaum, J. B. (2015). How, whether, why: Causal judgments as counterfactual contrasts. In D. C. Noelle, R. Dale, A. S. Warlaumont, J. Yoshimi, T. Matlock, C. D. Jennings, & P. P. Maglio (Eds.), *Proceedings of the 37th Annual Conference of the Cognitive Science Society* (pp. 782–787). Austin, TX: Cognitive Science Society.
- Gerstenberg, T., Peterson, M. F., Goodman, N. D., Lagnado, D. A., & Tenenbaum, J. B. (2017). Eye-tracking causality. *Psychological Science*, *28*, 1731–1744. <http://dx.doi.org/10.1177/0956797617713053>
- Gerstenberg, T., & Tenenbaum, J. B. (2016). Understanding "almost": Empirical and computational studies of near misses. In A. Papafragou, D. Grodner, D. Mirman, & J. C. Trueswell (Eds.), *Proceedings of the 38th Annual Conference of the Cognitive Science Society* (pp. 2777–2782). Austin, TX: Cognitive Science Society.
- Gerstenberg, T., & Tenenbaum, J. B. (2017). Intuitive theories. In M. Waldmann (Ed.), *Oxford handbook of causal reasoning* (pp. 515–548). New York, NY: Oxford University Press.
- Gerstenberg, T., Ullman, T. D., Nagel, J., Kleiman-Weiner, M., Lagnado, D. A., & Tenenbaum, J. B. (2018). Lucky or clever? From expectations to responsibility judgments. *Cognition*, *177*, 122–141. <http://dx.doi.org/10.1016/j.cognition.2018.03.019>
- Grice, H. P. (1975). Logic and conversation. In P. Cole & J. L. Morgan (Eds.), *Syntax and semantics: Vol. 3. Speech acts* (pp. 41–58). New York, NY: Wiley.
- Gureckis, T. M., Martin, J., McDonnell, J., Rich, A. S., Markant, D., Coenen, A., . . . Chan, P. (2016). psiTurk: An open-source framework for conducting replicable behavioral experiments online. *Behavior Research Methods*, *48*, 829–842. <http://dx.doi.org/10.3758/s13428-015-0642-8>

- Hall, N. (2007). Structural equations and causation. *Philosophical Studies*, 132, 109–136. <http://dx.doi.org/10.1007/s11098-006-9057-9>
- Halpern, J. Y. (2016). *Actual causality*. <http://dx.doi.org/10.7551/mitpress/10809.001.0001>
- Halpern, J. Y., & Hitchcock, C. (2015). Graded causation and defaults. *British Journal for the Philosophy of Science*, 66, 413–457. <http://dx.doi.org/10.1093/bjps/axt050>
- Harinen, T. (2017). Normal causes for normal effects: Reinvigorating the correspondence hypothesis about judgments of actual causation. *Erkenntnis*, 82, 1299–1320. <http://dx.doi.org/10.1007/s10670-017-9876-4>
- Hart, H. L. A., & Honoré, T. (1959/1985). *Causation in the law*. New York, NY: Oxford University Press.
- Henne, P., Niemi, L., Pinillos, Á., De Brigard, F., & Knobe, J. (2019). A counterfactual explanation for the action effect in causal judgment. *Cognition*, 190, 157–164. <http://dx.doi.org/10.1016/j.cognition.2019.05.006>
- Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, 15, 534–539. <http://dx.doi.org/10.1111/j.0956-7976.2004.00715.x>
- Hesslow, G. (1988). The problem of causal selection. In D. J. Hilton (Ed.), *Contemporary science and natural explanation: Commonsense conceptions of causality* (pp. 11–32). Brighton, United Kingdom: Harvester Press.
- Hilton, D. J. (1990). Conversational processes and causal explanation. *Psychological Bulletin*, 107, 65–81. <http://dx.doi.org/10.1037/0033-2909.107.1.65>
- Hilton, D. J., & Slugoski, B. R. (1986). Knowledge-based causal attribution: The abnormal conditions focus model. *Psychological Review*, 93, 75–88. <http://dx.doi.org/10.1037/0033-295X.93.1.75>
- Hitchcock, C. (2012). Portable causal dependence: A tale of consilience. *Philosophy of Science*, 79, 942–951. <http://dx.doi.org/10.1086/667899>
- Hitchcock, C., & Knobe, J. (2009). Cause and norm. *Journal of Philosophy*, 106, 587–612. <http://dx.doi.org/10.5840/jphil20091061128>
- Icard, T. F., Kominsky, J. F., & Knobe, J. (2017). Normality and actual causal strength. *Cognition*, 161, 80–93. <http://dx.doi.org/10.1016/j.cognition.2017.01.010>
- Johnson, S. G., & Rips, L. J. (2015). Do the right thing: The assumption of optimality in lay decision theory and causal judgment. *Cognitive Psychology*, 77, 42–76. <http://dx.doi.org/10.1016/j.cogpsych.2015.01.003>
- Kahneman, D., & Miller, D. T. (1986). Norm theory: Comparing reality to its alternatives. *Psychological Review*, 93, 136–153. <http://dx.doi.org/10.1037/0033-295X.93.2.136>
- Kahneman, D., & Tversky, A. (1982). The simulation heuristic. In D. Kahneman & A. Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases* (pp. 201–208). <http://dx.doi.org/10.1017/CBO9780511809477.015>
- Kennedy, A. (2015). Obergefell et al. v. Hodges, Director, Ohio Department of Health, et al. Congress. Retrieved from https://www.supremecourt.gov/opinions/14pdf/14-556_3204.pdf
- Kirfel, L., & Lagnado, D. A. (2018). Statistical norm effects in causal cognition. In T. T. Rogers, M. Rau, X. Zhu, & C. W. Kalish (Eds.), *Proceedings of the 40th Annual Conference of the Cognitive Science Society* (pp. 615–620). Austin, TX: Cognitive Science Society.
- Kirfel, L., & Lagnado, D. A. (2019). I know what you did last summer (and how often): Epistemic states and statistical normality in causal judgments. In A. K. Goel, C. M. Seifert, & C. Freksa (Eds.), *Proceedings of the 41st Annual Conference of the Cognitive Science Society* (pp. 575–581). Montreal, QB: Cognitive Science Society.
- Knobe, J., & Fraser, B. (2008). Causal judgment and moral judgment: Two experiments. In W. Sinnott-Armstrong (Ed.), *Moral psychology: Vol. 2. The cognitive science of morality: Intuition and diversity* (pp. 441–447). Cambridge, MA: MIT Press.
- Kominsky, J. F., Phillips, J., Gerstenberg, T., Lagnado, D., & Knobe, J. (2015). Causal superseding. *Cognition*, 137, 196–209. <http://dx.doi.org/10.1016/j.cognition.2015.01.013>
- Lagnado, D. A., Gerstenberg, T., & Zultan, R. (2013). Causal responsibility and counterfactuals. *Cognitive Science*, 37, 1036–1073. <http://dx.doi.org/10.1111/cogs.12054>
- Lombrozo, T. (2010). Causal-explanatory pluralism: How intentions, functions, and mechanisms influence causal ascriptions. *Cognitive Psychology*, 61, 303–332. <http://dx.doi.org/10.1016/j.cogpsych.2010.05.002>
- McGill, A. L., & Tenbrunsel, A. E. (2000). Mutability and propensity in causal selection. *Journal of Personality and Social Psychology*, 79, 677–689. <http://dx.doi.org/10.1037/0022-3514.79.5.677>
- McGrath, S. (2005). Causation by omission: A dilemma. *Philosophical Studies*, 123, 125–148. <http://dx.doi.org/10.1007/s11098-004-5216-z>
- Morris, A., Phillips, J., Gerstenberg, T., & Cushman, F. (2019). Quantitative causal selection patterns in token causation. *PLoS ONE*, 14(8), e0219704. <http://dx.doi.org/10.31234/osf.io/nq53z>
- Morris, A., Phillips, J., Icard, T., Knobe, J., Gerstenberg, T., & Cushman, F. (2018, April 26). *Judgments of actual causation approximate the effectiveness of interventions*. <http://dx.doi.org/10.31234/osf.io/nq53z>
- Pearl, J. (1999). Probabilities of causation: Three counterfactual interpretations and their identification. *Synthese*, 121, 93–149. <http://dx.doi.org/10.1023/A:1005233831499>
- Petrocelli, J. V., Percy, E. J., Sherman, S. J., & Tormala, Z. L. (2011). Counterfactual potency. *Journal of Personality and Social Psychology*, 100, 30–46. <http://dx.doi.org/10.1037/a0021523>
- Phillips, J., Luguri, J. B., & Knobe, J. (2015). Unifying morality’s influence on non-moral judgments: The relevance of alternative possibilities. *Cognition*, 145, 30–42. <http://dx.doi.org/10.1016/j.cognition.2015.08.001>
- Samland, J., Josephs, M., Waldmann, M. R., & Rakoczy, H. (2016). The role of prescriptive norms and knowledge in children’s and adults’ causal selection. *Journal of Experimental Psychology: General*, 145, 125–130. <http://dx.doi.org/10.1037/xge0000138>
- Samland, J., & Waldmann, M. R. (2016). How prescriptive norms influence causal inferences. *Cognition*, 156, 164–176. <http://dx.doi.org/10.1016/j.cognition.2016.07.007>
- Schaffer, J. (2005). Contrastive causation. *Philosophical Review*, 114, 327–358. <http://dx.doi.org/10.1215/00318108-114-3-327>
- Scholl, B. J., & Tremoulet, P. D. (2000). Perceptual causality and animacy. *Trends in Cognitive Sciences*, 4, 299–309. [http://dx.doi.org/10.1016/S1364-6613\(00\)01506-0](http://dx.doi.org/10.1016/S1364-6613(00)01506-0)
- Spranca, M., Minsk, E., & Baron, J. (1991). Omission and commission in judgment and choice. *Journal of Experimental Social Psychology*, 27, 76–105. [http://dx.doi.org/10.1016/0022-1031\(91\)90011-T](http://dx.doi.org/10.1016/0022-1031(91)90011-T)
- Stapleton, J. (2001). Legal cause: Cause-in-fact and the scope of liability for consequences. *Vanderbilt Law Review*, 54, 941–1009.
- Stapleton, J. (2008). Choosing what we mean by “causation” in the law. *Missouri Law Review*, 73, 433–480.
- Stephan, S., Willemsen, P., & Gerstenberg, T. (2017). Marbles in inaction: Counterfactual simulation and causation by omission. In G. Gunzelmann, A. Howes, T. Tenbrink, & E. Davelaar (Eds.), *Proceedings of the 39th Annual Conference of the Cognitive Science Society* (pp. 1132–1137). Austin, TX: Cognitive Science Society.
- Summers, A. (2018). Common-sense causation in the law. *Oxford Journal of Legal Studies*, 38, 793–821. <http://dx.doi.org/10.1093/ojls/gqy028>
- Sytsma, J., Livengood, J., & Rose, D. (2012). Two types of typicality: Rethinking the role of statistical typicality in ordinary causal attributions. *Studies in History and Philosophy of Biological and Biomedical Sciences*, 43, 814–820. <http://dx.doi.org/10.1016/j.shpsc.2012.05.009>
- Willemsen, P. (2016). Omissions and expectations: A new approach to the things we failed to do. *Synthese*, 195, 1–2. <http://dx.doi.org/10.1007/s11229-016-1284-9>
- Wolff, P. (2007). Representing causation. *Journal of Experimental Psychology: General*, 136, 82–111. <http://dx.doi.org/10.1037/0096-3445.136.1.82>

- Wolff, P., Barbey, A. K., & Hausknecht, M. (2010). For want of a nail: How absences cause events. *Journal of Experimental Psychology: General*, *139*, 191–221. <http://dx.doi.org/10.1037/a0018129>
- Wolff, P., Hausknecht, M., & Holmes, K. (2011). Absent causes, present effects: How omissions cause events. In J. Bohnemeyer & E. Pederson (Eds.), *Event representation in language and cognition* (pp. 228–252). New York, NY: Cambridge University Press.
- Woodward, J. (2006). Sensitive and insensitive causation. *Philosophical Review*, *115*, 1–50. <http://dx.doi.org/10.1215/00318108-2005-001>

Received May 2, 2019

Revision received June 28, 2019

Accepted July 11, 2019 ■