

Norms and the meaning of omissive enabling conditions

Paul Henne¹, Paul Bello², Sangeet Khemlani², and Felipe De Brigard¹

paul.henne@duke.edu, paul.bello@nrl.navy.mil, sangeet.khemlani@nrl.navy.mil, felipe.debrigard@duke.edu

¹Department of Philosophy, Duke University, Durham, NC 27708 USA

²Navy Center for Applied Research in Artificial Intelligence, Naval Research Laboratory, Washington, DC 20375 USA

Abstract

People often reason about omissions. One line of research shows that people can distinguish between the semantics of omissive causes and omissive enabling conditions: for instance, not flunking out of college enabled you (but didn't cause you) to graduate. Another line of work shows that people rely on the normative status of omissive events in inferring their causal role: if the outcome came about because the omission violated some norm, reasoners are more likely to select that omission as a cause. We designed a novel paradigm that tests how norms interact with the semantics of omissive enabling conditions. The paradigm concerns the circuitry of a mechanical device that plays music. Two experiments used the paradigm to stipulate norms and present a distinct set of possibilities to participants. Participants chose which causal verb best described the operations of the machine. The studies revealed that participants' responses are best predicted by their tendency to consider the semantics of omissive relations. In contrast, norms had little to no effect in participants' responses. We conclude by marshaling the evidence and considering what role norms may play in people's understanding of omissions.

Keywords: omissive causes; enabling; allowing; modal semantics; norms; mental models

Introduction

A railway gatekeeper's job is to open and close a crossing gate that lets trains pass. In 1902, the gatekeeper for the Somerset and Dorset railway was found guilty of manslaughter because he failed to close the gate (R. v. Pittwood, 1902). While he was at lunch, a train passed through the open gate and crashed into a horse and cart, killing one man and injuring another. The case describes an omissive cause: the jury held that the gatekeeper's failure to close the gate caused the death of an innocent bystander.

Omissive causation is a controversial topic amongst philosophers, psychologists, and legal scholars (Moore, 2009; Bernstein, 2015; Henne, Pinillos, & De Brigard, 2017). People have little difficulty in distinguishing which event was causal from alternative events that are non-causal. But causes are often easy to establish when they occur in a particular place and at a particular time; for instance, throwing a switch at a particular time causes the particular gate to close, so it is easy to identify the intervening action as the cause. Since omissive causes—absences, failures to act, scarcities, etc.—do not occur in any spatial or temporal frame, they present unique difficulties for causal reasoning and theories of causation.

One proposal suggests that norm violations affect causal judgments and play a fundamental role in establishing what constitutes a cause (Hitchcock & Knobe, 2009; McGrath, 2005; Hart & Honoré, 1985). In the railway example, the

gatekeeper was charged and found guilty because his occupation made it his responsibility to monitor the track. It may seem trivial that many other individuals—for instance, some passerby—also failed to close the gate, but previous philosophical treatments have difficulty explaining why only certain omissive causes are deemed relevant and not others (see McGrath, 2005; Bernstein, 2015). On the norm-based account, the passerby, unlike the gatekeeper, is not considered a cause, as there was no normative expectation for him to close the gate. The norm-based account provides an explanation for why people focus their attention on potential causes. Consistent with this view, recent studies show that reasoners view norm-violating omissions as causes but norm-preserving events as non-causes or as enablers (Henne et al., 2017; see also Clarke et al., 2013).

Nevertheless, some theorists question whether norms determine the meaning of omissive causal statements or whether norms simply bias causal judgments (Bernstein, 2014; 2017, p. 89-90). Consider the following statement:

1. The drought caused the famine.

Some argue that omissive causal statements as in (1) do not involve norms in any way, yet they are easy to comprehend. If norms were a central part of the meanings of causal relations, then the absence of any norm should render (1) uninterpretable (Bernstein, 2017). On such a view, norms may be relevant in *establishing* causal relations—such as in the train example—but they are not central to their meaning.

One clue for what it means to be a cause comes from the application of causal verbs: “causes,” “enables,” and “prevents.” Each verb refers to a relation between two events, and those relations have stark differences in their semantics. Psychological accounts of causal reasoning identify differences in the way people understand causal verbs (e.g., Goldvarg & Johnson-Laird, 2001; Sloman, Barbey, & Hotaling, 2009; Wolff, 2007). Accordingly, a viable theory of how people understand and infer omissive relations must distinguish the semantics between them. Consider the following two statements:

- 2a. An absence of light *causes* a flower to die.
- b. An absence of light *enables* a flower to die.

(2a) seems sensible, but (2b) does not, because (2b) implies that the flower can live without light. Likewise, in the following two statements:

- 3a. A lack of insecticides *causes* insects to thrive.
- b. A lack of insecticides *enables* insects to thrive.

(3b) seems sensible, but (3a) does not, because (3a) inappropriately guarantees that insects will thrive once insecticides are eliminated. The distinctions may be compelling, but until recently, no theory of causal reasoning could explain them.

A recent theory of omissive causation differentiates omissive causes from omissive enablers (Khemlani, Wasylyshyn, Briggs, & Bello, 2018). The theory is based on the idea that people represent causal scenarios by constructing and manipulating a set of discrete possibilities, i.e., mental models (Goldvarg & Johnson-Laird, 2001). The model-based theory—the “model theory,” for short—posits that omissive causes and omissive enabling conditions differ in the sets of possibilities to which they refer. On tasks that require reasoners to distinguish between the different relations, they should base their judgments on the semantics stipulated by the model theory (Khemlani et al., 2018). In contrast, if norms are central to the meaning of omissive relations, reasoners should base their decisions on norm-violations (Henne et al., 2017).

In what follows, we first delineate the predictions of the model theory and the norms hypotheses. We then describe an experimental paradigm that can test between the two predictions, and we present two novel experiments that test the competing predictions. The studies showed that reasoners separated omissive causes from omissive enabling conditions in a manner predicted by the model theory, and norm-violations had little effect on their behavior.

The model theory of causal reasoning

The mental model theory of human reasoning proposes that humans reason based on representing sets of possibilities (Johnson-Laird, 2006). The meanings of spatial relations, temporal relations, and causal relations refer to the sets of possibilities consistent with each relation (Goodwin & Johnson-Laird, 2005; Khemlani, Barbey, & Johnson-Laird, 2014). The model theory posits two systems of reasoning: a fast, intuitive system of reasoning constructs a single initial possibility—the “mental model”—to represent one or more assertions. Reasoners can formulate inferences rapidly by scanning that initial possibility, but those inferences are prone to error, because causal relations can be consistent with several possibilities. Errors can be corrected through deliberation, which is a process by which reasoners iteratively construct and consider alternative possibilities.

Khemlani and colleagues recently extended the model theory to account for reasoning about omissive causation (Bello et al., 2017; Khemlani et al., 2018). Their account explains why people distinguish different omissive relations (Table 1). It appeals to the idea that people rapidly construct initial mental models, and then flesh out those initial models into “fully explicit” models. On this view, a mental model is a privileged, default possibility to which an omissive causal relation refers, whereas fully explicit models represent all the possibilities consistent with the modal semantics of the relation. The following diagram depicts the mental model of the omissive causal relation described in (2a):

¬ light	death
---------	-------

where ‘¬’ denotes the symbol for negation (Khemlani, Orenes, & Johnson-Laird, 2012). Here, the lack of light is interpreted as a negated event, and it arranges the two events in the same chronological order in which they would occur. Hence, the model represents a single iconic possibility. When reasoners deliberate, they can consider all of the possibilities that accord with the modal semantics of omissive causation (Table 1). They can accordingly build fully explicit models of (2a), which are depicted in this diagram:

¬ light	death
light	death
light	¬ death

where each row represents a separate possibility. The bolded row represents the mental model. The latter two possibilities show that if the flower receives light, it may die anyway (for some other reason), or it may not die at all. But the theory predicts that reasoners should be less likely to think of these latter two possibilities at the outset because most reasoners only construct and reason with the mental model.

The theory posits that the mental model of omissive enabling conditions is the same as the mental model of omissive causation. Hence, the model of (3b) above is:

¬ insecticide	thrive
---------------	--------

It predicts that reasoners who draw conclusions on the basis of mental models should often conflate the two assertions (e.g., Wolff, 2007; Frosch & Johnson-Laird, 2011). When reasoners distinguish between omissive causes and enabling conditions, they should do so on the basis of their modal semantics, i.e., on the fully explicit models of the relations. The fully explicit models of an omissive enabling condition are depicted in the following diagram:

¬ insecticide	thrive
¬ insecticide	¬ thrive
insecticide	¬ thrive

Unlike omissive causes, omissive enabling conditions are consistent with the possibility in which both the cause and the effect do not hold, i.e., the situation in which insecticides are administered and insects subsequently do not thrive. Omissive enabling relations typically prohibit the possibility in which the cause and the effect both hold (A and B), e.g., the insects thrive even when they are sprayed with insecticide. But in some situations, omissive enabling relations can take on a weaker meaning and permit that possibility, as in, “The failure to cut the grass enabled it to grow.” The statement permits the possibility in which the grass is cut and it grows anyway (Table 1).

The model theory accordingly makes the following general hypothesis about semantics of omissive relations:

Semantics hypothesis: On tasks that require reasoners to distinguish between alternative causal relations, they should discriminate between omissive causes and omissive enabling conditions on the basis of the possibilities unique to each relation.

In contrast, reasoners are often susceptible to norm violations that affect their causal judgments (Henne et al., 2017). Hence, norm-based accounts posit the following hypothesis:

Norms hypothesis: When norms are available, reasoners distinguish between causal relations by focusing on those candidate events that violate norms. Events that violate norms should be considered causes, whereas those that do not violate norms should be rejected as causes.

In the next section, we describe a novel paradigm developed to test between the two hypotheses.

A paradigm for testing semantics and norms

Many existing paradigms test the meanings of omissive causes, but they do not typically encourage reasoners to consider and track multiple possibilities that are thought to be essential to the meanings (Henne et al., 2016; Khemlani et al., 2018; Wolff, Barbey, & Hausknecht, 2010; cf., Bello et al., 2017; Experiment 2). To try to overcome this limitation, we developed a novel paradigm that could be used to investigate how people distinguish the meanings of causal relations.

The paradigm made use of diagrams akin to those shown in Table 2. The basic diagram depicts a machine with a speaker, a red battery, a safety switch (which appeared green or else black), a blue wire, and an unnamed yellow component. The diagram could vary in several ways in order to depict different possibilities. For instance, the speaker could be playing or not playing (depicted as a series of soundwaves or not); the blue wire could be connected to the red battery, or else not connected; and the safety is green when it is *on* and black when it is *off*. The safety switch allowed for the establishment of a norm: participants were taught that whenever the safety is on (colored green), the blue wire *is not supposed to touch* the battery. And when the safety is off (colored black), the blue wire *is supposed to touch* the battery. Hence, the machine could be depicted in 2 (sound or no sound) x 2 (wire connected or disconnected) x 2 (safety on or off) = 8 different configurations. Depending on the particular condition in the experiments, those various configurations either violated or preserved norms. And those various configurations were either compatible with a particular omissive causal relation or incompatible with it.

The paradigm allows to directly compare predictions made by both the semantics hypothesis and the norms hypothesis. Experiments 1 and 2 provided participants a single trial in which they received three different diagrams (Table 2). After studying three different diagrams, participants were given a sentence completion task that tested their understanding of the scenario depicted.

Experiment 1

Experiment 1 presented participants with three separate diagrams. Participants only received diagrams that were either possible or impossible given causal and enabling relations: diagrams depicting context-dependent contingencies were not used in the experiment (Tables 1 and 2). Half the participants saw a set of diagrams in which each diagram was compatible with the following omissive enabling relation:

4. The blue wire not touching the red battery *allows* the speaker to play music.

The other half saw a set of diagrams that were compatible with (4) and the following causal assertion (5):

5. The blue wire not touching the red battery *causes* the speaker to play music.

Because the experiment avoided context-dependent contingencies, the remaining diagrams compatible with (5) were also compatible with (4), so the set of diagrams were ambiguous: the model theory predicts that participants should treat them as depicting both omissive causes and omissive enabling conditions (Khemlani et al., 2018, Experiment 4).

Participants were given a sentence completion task in which they chose the causal verb (“causes” or “allows”) to complete the following sentence:

The blue wire not touching the battery _____ the speaker to play music.

Notably, for the purposes of the study, the verb “allows” was treated as equivalent to “enables”. The experiment accordingly tested the prediction of the semantics hypothesis that reasoners should select the causal verb that matched the possibilities depicted. Hence, the semantics hypothesis predicts that reasoners should select “allows” more often for the enabling condition than the ambiguous condition. Experiment 1 also tested the norms hypothesis. Half the conditions in the study concerned abnormal situations in which the safety was off, and the other half concerned normal situations in which the safety was on, and participants were instructed that the blue wire is not supposed to touch the battery when the safety was on, and that it was supposed to touch the battery when the safety was off. The norms hypothesis predicts that people should be sensitive to norm violations, i.e., they should be more likely to select the verb “causes” to fill in the sentence provided for abnormal omissions compared to normal ones. In turn, normal omissions should be judged to be involved in enabling relations. The semantic hypothesis, however, posits that participants’ responses should not vary as a function of whether the condition was abnormal or normal—only as a function of which set of possibilities participants considered.

Omissive relation	The four possible contingencies between A and B							
	$\neg A$	B	$\neg A$	$\neg B$	A	$\neg B$	A	B
The lack of A causes B.	Mental model		Impossible		Possible		Context-dependent	
The lack of A enables B.	Mental model		Possible		Possible		Context-dependent	

Table 1: The table outlines the semantics of omissive causes and omissive enabling conditions. The rows separate omissive causes and omissive enabling conditions. The cells in each column describe whether each contingency is possible given a particular omissive relation. The mental models are always possible. The bolded column denotes the contingency diagnostic of omissive enabling conditions.

Methods

Participants A total of 822 adults participated in this study on Amazon Mechanical Turk (AMT). Of these participants, 59 did not complete the study, and 26 were excluded for failing to pass two attention checks. Data were analyzed with the remaining 796 participants ($M_{age} = 34$, $SD = 11.0$, age range = [18-71], 43% females).

Design and procedure Participants were randomly assigned to one of four possible conditions in a 2 (enabling vs. ambiguous) $\times$ 2 (normal vs. abnormal) between-participants design. Participants were acquainted with the machine in Figure 1 and its various components. After viewing and responding to instructions, participants were presented with three diagrams of configurations of the machine. The diagrams appeared on the screen simultaneously. To check their comprehension of the machine and the three separate possibilities, they matched the possibilities with descriptions provided in a dropdown menu, i.e., they chose from the following options to describe each of the three diagrams: (1) “The blue wire touches the battery, and the speaker plays music,” (2) “The blue wire touches the battery, and the speaker does not play music,” (3) “The blue wire does not touch the battery, and the speaker plays music,” and (4) “The blue wire does not touch the battery, and the speaker does not play music.” The order in which the possibilities were presented was fixed. In the normal condition, the machine’s safety was green, so the blue wire was not supposed to touch the battery. The blue wire not touching the battery is normal. In the abnormal condition, the machine’s safety was black, so the blue wire is supposed to touch the battery. The blue wire not touching the battery is abnormal. Participants were explicitly instructed to attend to the color of the safety and what the blue wire was supposed to do. They were then asked to think back to their observations and then fill in the verb in the sentence: “The blue wire not touching the battery _____ the speaker to play music.” Participants could choose between the verb “causes” or “allows” from a drop-down menu, and they could not proceed until a choice was made.

Post-experimental questionnaire Participants filled out a post-experimental questionnaire that asked them if they had paid attention and if they had taken the survey multiple times. Participants who reported affirmatively on either question were excluded.

Results and discussion

Figure 1 shows the proportion of participants who chose “allows” as a function of whether the condition was normal or not and as a function of whether the diagrams were consistent with the semantics for omissive enabling conditions or else ambiguous. Participants chose “allows” more often for diagrams consistent with enabling conditions than for ambiguous diagrams (74% vs. 54%; Mann-Whitney test, $z = 5.65$, $p < .0001$, Cliff’s $\delta = .19$). Participants selected “allows” more often when the diagrams were presented in a normal rather than an abnormal context—although this result was not statistically significant (68% vs. 62%; Mann-Whitney test, $z = 1.72$, $p = .09$, Cliff’s $\delta = .06$). A follow-up generalized logistical mixed-model (GLMM) regression further revealed that the difference in selection between the two conditions was inconsistent with a significant effect ($\beta = .00$, $p = .97$), as was the interaction between the two conditions ($\beta = .48$, $p = .12$). Nevertheless, a planned comparison revealed that for ambiguous diagrams, participants selected “allows” more often when the diagrams were presented in a normal context rather than an abnormal context (61% vs. 49%, Mann-Whitney test, $z = 2.35$, $p = .02$,




Figure 1: Proportion of participants who chose “allows” instead of “causes” in Experiment 1 as a function of whether participants saw normal or abnormal devices, and as a function of whether the diagrams were consistent with omissive enabling conditions only or consistent with both omissive causes and omissive enabling conditions. Error bars indicate 95% confidence intervals.






Condition	Diagrams presented to participants			
	-wire music	-wire -music	wire -music	wire music
Ambiguous	 (shown twice)	Not shown to participants		Not shown to participants
Enabling				Not shown to participants

Table 2: The three diagrams presented to participants in enabling and the ambiguous conditions. These diagrams all depict the normal conditions, i.e., the safety is on, so the blue wire is not supposed to touch the red battery. In Experiment 1, participants were not provided with written cues about the condition of the safety switch. In Experiment 2, those cues were provided (as in the diagrams above).

Cliff's $\delta = .12$). All data and analysis code available at <https://osf.io/jf36w/>. The result provides some support for the norm hypothesis, which predicts that people should be more likely to select “causes” (and less likely to select “allows”) for abnormal contexts. It also suggests that participants were sensitive to the norm manipulation: they comprehended the norms and took them into account in making their selections. If they had not, they would have shown no sensitivity to whether the diagrams were in normal or abnormal context. Yet, an analogous comparison for diagrams consistent with omissive enabling conditions was not reliable, and so the study revealed mixed support for the norms hypothesis.

Experiment 1 corroborated the prediction that reasoners interpret omissive causes and enabling conditions in accordance with the semantics outlined by the model theory. Reasoners in the enabling condition selected “allows” more often than those in the ambiguous condition. Moreover, participants’ responses did not depend on whether a norm had been violated or not. If, as the norms hypothesis states, abnormal situations help reasoners choose which candidate events constitute causes, then those abnormalities appeared to have no effect on participants’ tendencies to select appropriate causal relations.

One limitation of Experiment 1 is that reasoners may have simply failed to recognize abnormalities in the first place, i.e., they may not have encoded the black safety switch’s color, which was designed to serve as a cue that a norm had been violated. Another limitation of the study is that participants evaluated only one set of three diagrams and only one causal relation. Experiment 2 corrected for both of these limitations.

Experiment 2

Because participants may not have picked up on the norm distinction between conditions in Experiment 1, we sought to ensure that the difference was salient in Experiment 2. Hence, rather than just identifying the color as the difference in norms, Experiment 2 added the verbal cues “SAFETY ON” and “SAFETY OFF” to the diagrams (see Table 2). Moreover, the

study employed a within-participants design to further validate the findings from Experiment 1 supporting the semantics hypothesis. Hence, each participant saw four distinct sets of three diagrams.

Methods

Participants A total of 215 adults participated in this study on AMT. Of these, 21 participants were excluded for failing to pass two attention checks. Data were analyzed with the remaining 194 participants ($M_{age} = 33.82$, $SD = 9.39$, age range [18-68] 40% females).




Figure 2: The stimuli used in Experiment 2. Participants in the abnormal condition saw the machine with the black safety i.e., in the off position (top diagram), and those in the normal condition saw the machine with the green safety, i.e., in the on position (bottom diagram).

Design and procedure Participants acted as their own controls and received all four possible conditions in a 2 (enabling vs. ambiguous) x 2 (normal vs. abnormal) within-participants design. As in Experiment 1, participants were acquainted with the machine and its various components. The model conditions were constructed just as they were in Experiment 1 (see Table 2). The stimuli were modified such that the normal and abnormal conditions were more salient by adding the words “SAFETY ON” and “SAFETY OFF” to the diagrams (Figure 2).

Results and discussion

Figure 3 shows the proportion of participants who chose “allows” as a function of whether the condition was normal or not and as a function of whether the diagrams were consistent with the semantics for omissive enabling conditions or else ambiguous. As in Experiment 1, participants selected “allows” more often when the diagrams depicted possibilities uniquely consistent with omissive allowing relations rather than ambiguous possibilities (77% vs. 53%; Mann-Whitney test, $z = 7.00$, $p < .0001$, Cliff’s $\delta = .24$). They didn’t reliably select “allows” more often for diagrams in an abnormal vs. a normal context (66% vs. 64%; Mann-Whitney test, $z = .83$, $p = .41$, Cliff’s $\delta = .03$). GLMM regression analyses likewise corroborated the nonparametric analyses: it yielded an effect of whether the diagrams were ambiguous or consistent with omissive allowing conditions ($B = 1.14$, $p < .0001$), but no effect of normality ($B = .15$, $p = .54$) and no interaction ($B = .02$, $p = .95$).

A planned comparison revealed that for ambiguous diagrams, participants did not reliably select “allows” more often for abnormal than normal contexts (55% vs. 52%;




Figure 3: Proportion of participants who chose “allows” instead of “causes” in Experiment 2 as a function of whether participants saw normal or abnormal devices, and as a function of whether the diagrams were consistent with omissive enabling conditions only or ambiguous. Error bars indicate 95% confidence intervals.

Mann-Whitney test, $z = .61$, $p = .54$, Cliff’s $\delta = .03$). Hence, overall, the results are not consistent with the norms hypothesis, which states that reasoners should be more likely to select “causes” (and less likely to select “allows”) for abnormal vs. normal contexts.

Just as in Experiment 1, the results of Experiment 2 support the prediction of the model theory with respect to its predictions about the semantics of omissive enabling conditions. Reasoners in the enabling condition selected “allows” more often than those in the ambiguous condition. Moreover, participants’ responses did not depend on whether a norm had been violated or not.

General Discussion

Two experiments were designed to test how participants judge the causal effect omissive events have on outcomes. The experiments corroborated a recent theory of omissive causation, which predicts that reasoners should be able to distinguish omissive enabling conditions from other sorts of omissive relation (Khemlani et al., 2018). Moreover, the results showed that norm violations cannot explain the semantic difference between causes and enabling conditions.

The results of these studies can help refine the role that norms play in causal reasoning. As Henne and colleagues (2017) show, norms help select potential causes and distinguish them from irrelevant non-causes. The present studies, however, show that norms do not always explain the difference between causes and enablers. When reasoners consider the distinctive possibilities consistent with enabling conditions norms have little to no effect on causal judgment. When reasoners consider only the mental model, norms may have a more prominent effect on causal judgment. A more robust extension of the model theory, i.e., one that explains how norms are represented and how they modulate possibilities could potentially explain both the semantics between different causal verbs as well as how reasoners isolate potential causes from non-causes. Such a theory would also have to be contrasted with recent models of causal strength that could potentially explain the norm effects and the results predicted by the model theory (Icard, Kominsky, & Knobe, 2018).

Acknowledgments

This project was supported by an Office of Naval Research award (N00014-17-1-2603) to Felipe De Brigard.

References

- Bello, P., Wasylyshyn, C., Briggs, G., & Khemlani, S. (2017). Contrasts in reasoning about omissions. In G. Gunzelmann, A. Howes, T. Tenbrink, & E. Davelaar (Eds.), *Proceedings of the 39th Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.
- Bernstein, S. (2014). Omissions as possibilities. *Philosophical Studies*, 167.
- Bernstein, S. (2015). The metaphysics of omissions. *Philosophy Compass*, 10, 208-218.

- Bernstein, S. (2017). Intuitions and the metaphysics of causation. *Experimental Metaphysics*, 75.
- Clarke, R. et al. (2013). Causation, norms, and omissions: A study of causal judgments. *Philosophical Psychology*, 28, 279-93.
- Frosch, C.A., & Johnson-Laird, P.N. (2011). Is everyday causation deterministic or probabilistic? *Acta Psychologica*, 137, 280-291.
- Goldvarg, E., & Johnson-Laird, P. (2001). Naïve causality: A mental model theory of causal meaning and reasoning. *Cognitive Science*, 25, 565-610.
- Goodwin, G.P., & Johnson-Laird, P.N. (2005). Reasoning about relations. *Psychological Review*, 112, 468-493.
- Halpern, J.Y., & Hitchcock, C. (2014). Graded causation and defaults. *British Journal for the Philosophy of Science*, 66, 413-457.
- Hart, H.L.A., & Honoré, T. (1985). *Causation in the Law*, 2nd ed., Oxford: Clarendon.
- Henne, P., Pinillos, Á., & De Brigard, F. (2017). Cause by omission and norm: Not Watering Plants. *Australasian Journal of Philosophy*.
- Icard, T. F., Kominsky, J. F., & Knobe, J. (2017). Normality and actual causal strength. *Cognition*, 161, 80-93.
- Johnson-Laird, P.N. (2006). *How we reason*. NY: Oxford University Press.
- Johnson-Laird, P. N., & Khemlani, S. (2017). Mental models and causation. In M. Waldmann (Ed.), *Oxford Handbook of Causal Reasoning*. Oxford: Oxford University Press.
- Johnson-Laird, P. N., Khemlani, S., & Goodwin, G.P. (2015). Logic, probability, and human reasoning. *Trends in Cognitive Sciences*, 19.
- Khemlani, S., Barbey, A., & Johnson-Laird, P. N. (2014). Causal reasoning with mental models. *Frontiers in Human Neuroscience*, 8, 849, 1-15.
- Khemlani, S., Orenes, I., & Johnson-Laird, P.N. (2012). Negation: a theory of its meaning, representation, and use. *Journal of Cognitive Psychology*, 24.
- Khemlani, S., Wasylyshyn, C., Briggs, G., & Bello, P. (2018). Mental models and omissive causation. *Memory & Cognition*, 46, 1344-1359.
- McGrath, S. (2005). Causation by omission: A dilemma. *Philosophical Studies*, 123.
- Wolff, P. (2007). Representing causation. *Journal of Experimental Psychology: General*, 136, 82-111.
- Wolff, P., Barbey, A., & Hausknecht, A. (2010). For want of a nail: How absences cause events. *Journal of Experimental Psychology: General*, 139.