

Contrasts in reasoning about omissions

Paul Bello, Christina Wasylyshyn, Gordon Briggs, Sangeet Khemlani

{paul.bello, christina.wasylyshyn, gordon.briggs.ctr, sangeet.khemlani}@nrl.navy.mil

Naval Research Laboratory, Overlook Ave. S.W.

Washington, DC 20375 USA

Abstract

Omissions figure prominently in causal reasoning from diagnosis to ascriptions of negligence. One philosophical proposal posits that omissions are accompanied by a contrasting alternative that describes a case of orthodox (non-omissive) causation (Schaffer, 2005; Bernstein, 2014). A psychological hypothesis can be drawn from this contrast view of omissions: by default, humans should interpret omissive causations as representing at least two possibilities, i.e., a possibility representing the omission and a possibility representing a contrast. The theory of mental models supposes that reasoners construct only one possibility (the omission) by default, and that they consider separate alternative possibilities in sequential order. Two experiments test the contrast hypothesis against the model theory, and find evidence in favor of the model-theoretic account.

Keywords: omissive causation; mental models; reasoning; contrasts

Introduction

A mechanical failure in a car to start causes a missed meeting. A friend's broken promise causes hurt feelings. The lack of rainfall causes drought. Each is a case of omissive causation, where an omission or lack of some event brings about some effect. Take the following oft-cited example:

You come home after a business trip to find your rosebushes desiccated and ruined. You learn from your neighbor that your gardener did not show up to water the plants.

Omissive causes feature in both prediction and explanation. Intuition suggests that the death of the roses is explained by the failure of the gardener to show up. Similarly, we would expect that future failures of this kind would yield the same result. And omissions are often invoked in moral judgment. If you had signed a contract with the gardener and were especially litigious, you would have grounds to sue for damages. While we take for granted the fact that omissions are causes, omissions pose deep puzzles for theorists who wish to treat them in much the same way that "orthodox" (i.e., non-omissive) causes are treated.

It seems reasonable to think that orthodox causation concerns relations between events. But omissions are non-events, and it is unclear how a non-event can be an argument to a causal relation. One idea is that omissions could be nothing at all (Clarke, 2014; Beebe, 2004), but this notion fails to explain why omissions seem to serve as sensible causal agents, as in, e.g., *the lack of medicine caused sickness*. If omissions were nothing, then they

couldn't be thought of as causes. Another proposal is that omissions denote a non-actualized possibility (Bernstein, 2014), such as that of the gardener showing up to water the bushes. Bernstein invokes the machinery of possible worlds to argue that omissions involve "counterpart relations" between actual omitted events and non-actualized contrast events at close-by possible worlds. A related idea in Schaffer (2005) is that omissions represent actual events, e.g., the event that occurred instead of the gardener showing up. The shared assumption from these latter two proposals about the nature of omissions is that they are definable in terms of contrasts between events (Schaffer, 2005), or between the omitted event and a non-actualized possibility (Bernstein, 2014).

Bernstein and Schaffer's accounts, while distinct, share the surface similarity of basing omissions on contrasts. In both cases, each theorist argues that an adequate metaphysics of omissions ought to not run afoul of human intuitions. In this spirit of consilience between intuition and metaphysical theorizing, it is worth exploring whether or not contrasts are present in mental representations of omissive causes and the inferences reasoners draw from them. The present paper explores if and how statements about omissions automatically refer to representations of contrasting events. Evidence from the psychology of counterfactual reasoning suggests that reasoners are in principle capable of maintaining two separate possibilities (Byrne, 2005), but whether they do so for reasoning about omissions remains unknown. If a strict interpretation of the contrast view is correct, reasoners should interpret omissive causation as referring to two representations by default.

In what follows, we review a psychological account of omissive-causal reasoning with mental models (Bello & Khemlani, 2015). The theory predicts that reasoners tend to interpret omissive causation as referring to a single possibility: one in which the omitted event happens (e.g., the gardener fails to water the flowers), and the result follows (e.g., the flowers die). The theory posits that reasoners can potentially think about other possibilities including contrasts, e.g., the situation in which the gardener waters the flowers and they don't die, or the situation in which the gardener waters the flowers and they die for some other reason. But these alternative possibilities demand additional effort, and so by initially considering only one (non-contrasting) possibility, reasoners reduce the load on their working memories.

We describe two studies that test between the different hypotheses. The experiments support the model-based account in which reasoners do not represent contrasting possibilities by default, but instead consider alternatives

sequentially and in a systematic order. The paper concludes by discussing other puzzling aspects of omissive causation and plans for future research.

The model theory

The mental model theory of reasoning – or “model theory” for short – posits that reasoners draw conclusions by building and scanning *mental models*, or iconic representations of possibilities (Johnson-Laird, 2006; Johnson-Laird & Byrne, 1991, Goldvarg & Johnson-Laird, 2001; Goodwin & Johnson-Laird 2005). The model theory makes three central assumptions:

1. **The principle of iconic possibilities.** The contents of perception, memory, language, or imagination yield models, i.e., sets of discrete possibilities. Models are iconic, i.e., they are isomorphic to the structure of what they represent (Peirce, 1931-1958, Vol. 4), but they can also contain abstract tokens, such as a symbol denoting negation (Khemlani, Orenes, & Johnson-Laird, 2012). And they can represent temporal sequences of events as discrete possibilities that unfold in time the way events do (Khemlani & Johnson-Laird, 2013).
2. **The principle of parsimony.** Models require maintenance in working memory, and so inferences that demand more models are more difficult and take longer than those that demand fewer models. Hence, the theory posits two primary systems for reasoning: a fast system builds and scans models without the use of working memory, and so it posits that reasoners tend to reason with a single *mental model* in most scenarios. A slower system revises and rebuilds models, and it searches for alternative models consistent with the premises. It can correct the errors and biases that the fast system yields, but it is subject to the limitations of working memory.
3. **The principle of truth.** Reasoners initially build models that represent only what is true in a compound clause, and not what is false. They can flesh out the initial mental models to yield a set of *fully-explicit models*, i.e., those possibilities that denote both true, false, possible, and impossible scenarios. Fully-explicit models form a complete representation of the possibilities to which a statement refers.

To illustrate these three principles, we now turn to summarizing the theory of omissive causation presented in (Bello & Khemlani, 2015).

A model-based account of omissive causation

According to the model theory, different sorts of causal verbs refer to different sets of mental models (Goldvarg & Johnson-Laird, 2001; Khemlani, Barbey, & Johnson-Laird, 2014). For instance, the statement, *acid causes flowers to die*, refers to three separate possibilities that constitute a fully-explicit model, which can be depicted in the following diagram:

acid	death
– acid	– death
– acid	death

Each row of the diagram represents a different temporally-ordered possibility that renders the statement true. Hence, the first row denotes the possibility in which acid is introduced and the flowers die; and the latter two rows denote possibilities in which acid isn’t introduced and the flowers do not die (row 2) or die anyway (row 3). The model does not represent situations inconsistent with the statement (e.g., the situation in which acid is introduced and the flowers do not die, or any situation in which death occurs before acid is introduced). Moreover, maintaining three separate possibilities is difficult for reasoners, and the principle of parsimony implies that most reasoners only construct the first possibility, i.e., the mental model:

acid	death
------	-------

The mental model can be scanned and combined with other premises to yield inferences rapidly, but reasoners who rely on the mental model alone are prone to make reasoning errors on certain inferences. Moreover, each additional model in the set of fully-explicit models above demands working memory resources, and reasoners should be progressively less likely to consider them.

Omissive causation operates similarly to orthodox causation under the model theory, with the proviso that omissions imply that the antecedent events are negated (Bello & Khemlani, 2015) and negations increase difficulty (Khemlani et al., 2012). For instance, the statement, *the lack of water causes flowers to die*, refers to the following mental model:

– water	death
---------	-------

which can be fleshed out into the following fully-explicit models:

– water	death
water	– death
water	death

And so, just as in the case of orthodox causation, the model theory predicts that reasoners should often build only the mental model (i.e., the possibility in which there is a lack of water and the flowers die). Those who consider additional possibilities should construct the second possibility less often than the first, and the third possibility less often than the second.

Models and contrasts

Do omissive causes entail contrasts? If so, then a central assumption of the model theory would be incorrect. That is, a default representation of contrasts would imply that statements such as *the lack of water causes flowers to die* should refer to the following two models:

– water	death
water	– death

instead of just one mental model (see above). Reasoners do appear to consider the contrasting possibility often. In recent studies by Briggs and colleagues, participants evaluated

omissive causal relations (of a structure akin to: *the lack of A causes B*) by assessing whether four separate scenarios (*not-A and B*, *A and not-B*, *A and B*, and *not-A and not-B*) were possible given the truth of the relation. Participants in one study, for instance, selected *not-A and B* at ceiling, and they selected *A and not-B* close to ceiling (98% and 83%, respectively; see Briggs et al., under review, Experiment 3). The preponderance of *A and not-B* responses lends some tentative support to the idea that omissions are understood in terms of contrasts, as some metaphysicians have suggested. But, the data are also consistent with the view that reasoners select the contrasting possibility (*A and not-B*) only after considering the mental model (*not-A and B*) first. No studies directly test between the contrast view and the model theory, and so we carried out two experiments in which reasoners made inferences about omissive causation. Certain inferences should be less error-prone if reasoners represent contrasting possibilities, but Experiment 1 showed no such improvement. Experiment 2 showed that reasoners spontaneously generate contrasting possibilities less often – and after – they represent possibilities corresponding to mental models. Both studies support the predictions of the model theory.

Experiment 1

Experiment 1 tests reasoners’ inferences about the pattern of reasoning known as *modus tollens*, which is an inference in sentential logic of the following form:

If A then C.
 Not C.
 Therefore, not A.

The inference is *valid* because it is true in every case that the premises are true (Jeffrey, 1981, p. 1). But, reasoners have difficulty with *modus tollens* inferences: they tend to respond that nothing follows from the premises instead of inferring that *not A* follows (Nickerson, 2015, p. 41 et seq.). A causal version of the inference is as follows:

Overexposure to UV light causes snowblindness.
A particular mountaineer doesn’t have snowblindness.
What follows?

A valid conclusion from these premises is that *the mountaineer isn’t overexposed to UV light*. But, even in the causal domain, many reasoners have difficulty generating the valid conclusion, and they instead respond that nothing follows (Cummins, Lubart, Alksnis, & Rist, 1991). The model theory explains why: if reasoners represent only a single mental model of *overexposure to UV light causes snowblindness*, e.g.,

UV-overexposure snowblindness

then that single possibility does not correspond to the possibility referred to in the second premise: *a mountaineer doesn’t have snowblindness*. And so, reasoners respond that nothing follows. Only those reasoners who construct the fully explicit models of the causal relation:

UV-overexposure snowblindness
 – UV-overexposure – snowblindness
 – UV-overexposure snowblindness

can make the valid deductive inference, because the second premise corresponds to the second possibility above.

The same prediction, *mutatis mutandis*, holds for omissive causal relations. Consider the following inference:

A lack of vitamin C causes scurvy.
A particular sailor doesn’t have scurvy.
What follows?

If, as the model theory predicts, reasoners represent only a single mental model, e.g.,

– vitamin-C scurvy

then they should have difficulty drawing a valid conclusion from the premises. If, however, reasoners construct both the mental model and its contrasting possibility, e.g.,

– vitamin-C scurvy
 vitamin-C – scurvy

then they should be more likely to make the valid conclusion that *the sailor doesn’t have a vitamin C deficiency*. Hence, the contrastive view of omissive causation predicts that reasoners should respond more accurately on *modus tollens* inferences when they concern omissions than when they concern orthodox causation.

To test this prediction, participants in Experiment 1 wrote out their natural responses to short vignettes concerning omissive and orthodox causal reasoning arguments.

Method

Participants. Thirty participants volunteered through the Amazon Mechanical Turk online platform (see Paolacci, Chandler, & Ipeirotis, 2010, for a review). Fourteen participants reported no formal logic or advanced mathematical training and the remaining reported introductory to advanced training in logic. All participants were native English speakers.

Design, procedure, and materials. Participants carried out the experiment on a computer screen. The study was designed in psiTurk (Gureckis et al., 2015). After reading instructions, participants completed eight experimental problems. Half the problems concerned omissive causation by making use of the word “absence” to establish an omission; and the other half concerned orthodox causation by using the word “presence”. Each problem comprised two premises. The first premise always established the presence or absence of a causal relation (e.g., *A causes B*). For half of the problems, the second premise asserted that the event (*B*) occurred (and therefore allowed participants to draw an inference known as *affirming the consequent*), and for the other half, the premise asserted that the event did not occur (*not-B*), and so participants could draw a *modus tollens* inference. Participants wrote out responses to the question “What, if anything, follows?” An example problem is as follows:

Suppose the following statements are true:

1. The [presence/absence] of a particular part causes a machine to fail.
2. On a particular day, the machine [did/didn't] fail.
What, if anything, follows?

The information for each problem was presented simultaneously, and participants were prevented from continuing to the next problem until they typed in at least one possibility. Participants were informed that they should write out that “nothing followed” if they thought there was not enough information in the premises to make any conclusion with certainty. The materials were drawn from four domains: biology, nature, socioeconomics, and mechanics. The presentation order of the content and problem type of the vignettes was randomized.

Results and discussion

Two coders blind to the predictions of the study judged whether participants’ natural responses were accurate or inaccurate; they agreed on 99% of trials (Cohen’s $\kappa = .99$). Table 1 shows the percentage of accurate responses in Experiment 1 as a function of whether the inference concerned omissive or orthodox causation. Across the study as a whole, participants produced more accurate responses for orthodox causation than for omissive causation (47% vs. 32% correct; Wilcoxon test, $z = 2.49$, $p = .01$, Cliff’s $\delta = .15$), which is the opposite of the pattern predicted by the contrast view. As in previous studies (e.g., Cummins et al., 1991), participants produced more accurate responses for *modus tollens* inferences than for *affirming the consequent* inferences (63% vs. 15% correct; Wilcoxon test, $z = 6.65$, $p < .0001$, Cliff’s $\delta = .48$). The interaction between the type of causation and the type of inference was not reliable (Wilcoxon test, $z = .64$, $p = .52$, Cliff’s $\delta = .03$).

Participants in Experiment 1 violated the prediction of the contrast view: they were less accurate for inferences concerning omissive causation than for those concerning orthodox causation. Indeed, their patterns of inference corroborate the model theory, which predicts that inferences about omissive causes should be slightly more difficult because reasoners represent negated possibilities.

Inference	Type of causation	
	Orthodox <i>The presence of A...</i>	Omissive <i>The absence of A...</i>
<i>Affirming the consequent:</i> A causes B. B. What, if anything, follows?	22%	8%
<i>Modus tollens:</i> A causes B. Not B. What, if anything, follows?	72%	55%

Table 1. Proportion of correct responses in Experiment 1 as a function of the type of inference and the type of causation.

Experiment 2

Experiment 1 concerned inferences, and its results suggest that reasoners do not make use of contrasting possibilities in their *modus tollens* or *affirming the consequent* inferences. But, the data do not conclusively establish whether or not reasoners represent the contrast by default. After all, people might initially represent such possibilities but fail to consider them when drawing inferences. Hence, Experiment 2 tested reasoners’ interpretations of omissive causation directly. It elicited natural responses to the different possibilities for orthodox and omissive cause and enabling conditions. Participants read a single short premise and were asked to list the possibilities that correspond to each premise. We analyzed the order in which participants constructed each possibility, as well as the first possibility they constructed. The contrast view predicts that reasoners should construct the possibilities that correspond to *not-A and B* and *A and not-B* equally often when they interpret omissive causation. The model theory predicts that reasoners should construct the possibility that corresponds to *not-A and B* first, then (if at all) the possibility that corresponds to *A and not-B*, and finally (if at all) the possibility that corresponds to *A and B*. And the theory predicts an analogous trend in latencies: reasoners should build *not-A and B* faster than *A and not-B*, and they should build *A and not-B* faster than *A and B*.

Method

Participants. Thirty-one participants volunteered through the Amazon Mechanical Turk online platform. Twenty-two participants reported no formal logic or advanced mathematical training and the remaining reported introductory to advanced training in logic. All were native English speakers.

Design, procedure, and materials. Participants completed two practice problems and eight experimental problems, and they acted as their own controls. Each problem presented one premise that consisted of two events and a causal verb. The experiment manipulated whether the first event concerned orthodox or omissive causation: half the problems used the word “presence” and the other half used the word “absence.” The experiment also manipulated the relevant causal relation: half the problems concerned causation and half concerned enabling conditions, though for brevity we analyze only those problems concerning causation below. An example problem is as follows:

Suppose the following statement is true:

The [presence/absence] of a particular preservative [causes/enables] a substance to decay.

What is possible given the above statement?

Participants were then asked to construct a list of possibilities using pre-populated drop-down menus. Figure 1 shows an example of the interface used in Experiment 2. Participants could choose any combination of the possibilities from the drop-down menus, they could change




Figure 1. The interface used to elicit responses in Experiment 2. Participants completed sentences using drop-down menus and added possibilities using a button marked “+”.

their answer choices at will, and they could add additional sentences if they thought the statement was true in a number of possibilities; but, the interface allowed the construction of at most four different sentences. The presentation order of the trials was randomized. The order in which the participants endorsed possibilities was recorded, as was the latency between when the premises appeared and when participants pushed a button to finish the trial.

Results and discussion

Table 2 shows the percentage of trials on which participants constructed the four possible sentences as a function of whether the premise in the trial concerned an orthodox or an omissive causal relation. The table also shows, in parentheses, the percentages of trials on which a given sentence appeared first in the set of sentences constructed by the participants.

For omissive causation trials, participants constructed *not-A and B* more often than *A and not-B* (85% vs. 69%, respectively; Wilcoxon test, $z = 2.88$, $p = .003$, Cliff’s $\delta = .03$), in violation of the contrast view. Instead, the data corroborate the trend predicted by the model theory; participants constructed *not-A and B* most often (85% of trials), then *A and not-B* (69%), then *A & B* (47%), and rarely *not-A and not-B* (19%). A nonparametric trend test revealed a significant trend in their responses (Page’s trend test, $z = 5.16$, $p < .0001$).

One way of understanding participants’ performance is to examine only the first sentence in the set of sentences they constructed: doing so allows for a coarse analysis of their online preferences for possibilities. Participants constructed *not-A and B* as a first sentence more often than *A and not-B*

Type of causation	The four sentences (in abbreviated form)			
	<i>A & B</i>	<i>A & ¬B</i>	<i>¬A & B</i>	<i>¬A & ¬B</i>
Orthodox	100 (100)	6 (0)	31 (0)	74 (0)
Omissive	47 (26)	69 (26)	85 (48)	19 (0)

Table 2. Percentages of trials on which participants in Experiment 2 constructed four separate sentences for trials that concerned omissive and orthodox causal relations. “¬” denotes negation. In parentheses: percentages of trials on which a sentence appeared first in the set constructed by participants. (Not shown: data from trials that concerned enabling conditions.)

(48% vs. 26%, respectively; Wilcoxon test, $z = 2.06$, $p = .04$; Cliff’s $\delta = .22$). And they constructed *A and not-B* and *A and B* equally as often (26% vs. 26%). We recorded the latency between when the premises appeared to when participants made a response. While those latencies are inflated to include the amount of time they read the premises, they nevertheless revealed that participants were faster to construct *not-A and B* (26 s) as a first sentence compared to *A and not-B* (38 s); a test on their overall selections corroborated the trend predicted by the model theory (Jonckheere’s trend test, $z = 2.95$, $p = .001$).

Participants’ responses to orthodox causal relations likewise corroborated the predictions of the model theory. Every participant responded *A and B*, 74% of participants responded *not-A and not-B*, and 31% of participants responded *not-A and B* (Page’s trend test, $z = 3.54$, $p < .001$). And, every participant constructed *A and B* first.

The results largely corroborate the predictions of the model theory. The significant trends in the proportions of constructing the different sentences suggest that reasoners consider possibilities sequentially.

General discussion

Two experiments showed that when people interpret and reason about omissive causal relations, they do not represent contrasting alternatives by default. Reasoners can, in principle, hold two separate possibilities in mind – they seem to do precisely that when reasoning about counterfactual assertions (Byrne, 2005). But, as the present studies show, they tend to interpret omissive causes as referring to a single model of a negated cause and its associated effect (Experiment 1). When they are asked to list possibilities, they list the mental model earlier, more often, and faster than contrasting possibilities (Experiment 2). These results corroborate the model theory of omissive causation (Bello & Khemlani, 2015; Briggs et al., under review), and no alternative theory of omissive causation, whether psychological (Wolff et al., 2010) or philosophical, presently account for the results from the two studies.

The model theory provides a specific ordering on what kinds of possibilities reasoners consider. It is a process theory that explains why some possibilities (the mental models) are considered by default and why others (the alternative models) demand additional cognitive resources to construct. And it provides some constraints on the contents of alternative properties, though the specific contents depend on the semantics of the particular verbs used to describe omissive-causal premises.

How are contrast events identified in the first place? Both Bernstein and Schaffer highlight this open question, and they suggest that its answer is critical for shoring up their respective theories. Philosophers who interpret omissive causation using possible worlds are faced with developing a theory that specifies how to pick through the infinitude of possible worlds to find those that contain the most relevant contrast event. Actual-event theorists like Schaffer shoulder the same explanatory burden. To illustrate, it is perfectly

reasonable to think that the Queen of England might have shown up to water the rosebushes. This would have prevented their death, after all. Does her failure to do so qualify her as an omissive cause of their dying? What explains why the gardener's failure to do so is a better candidate for the actual cause of their death?

Recent work shows norms and pragmatics help establish relevant contrasts when reasoning about omissions (see Henne, Pinillos, & De Brigard, 2016). Analogously, the model theory posits that background knowledge of the meanings of words and their contexts can introduce relations and block the construction of certain possibilities in a process known as *modulation* (Johnson-Laird & Byrne, 2002). For instance, reasoners often infer a temporal relation from the conjunction, *Mary studied and she passed her test*, such that she studied *before* she passed her test (Juhos et al., 2012). In the case of omissive causes, modulation may rely on knowledge of norms to introduce relations or contents for the different contrasting possibilities, and future work will investigate the processes by which norms bias the representation of omissive causes.

Acknowledgments

This work was supported by grants from the Office of Naval Research to PB and SK. We are grateful to Knexus Research Corporation for their help in conducting the experiments, and we are especially grateful to Ruth Byrne, whose research on counterfactual inferences inspired the design of Experiment 1. Finally, we thank Monica Bucciarelli, Felipe de Brigard, Todd Gureckis, Paul Henne, Laura Hiatt, Phil Johnson-Laird, Greg Murphy, L.A. Paul, Bob Rehder, and Greg Trafton for their advice and comments.

References

Beebe, H. (2004). Causing and nothingness. In J. Collins, N. Hall, and L. A. Paul (Eds.), *Causation and Counterfactuals*, Cambridge, MA: The MIT Press.

Bello, P., & Khemlani, S. (2015). A model-based theory of omissive causation. In *Proceedings of the 37th Annual Meeting of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.

Bernstein, S. (2014). Omissions as possibilities. *Philosophical Studies*, 167.

Briggs, G., Wasylyshyn, C., Bello, P., & Khemlani, S. (under review). Mental models of omissive causation.

Byrne, R. (2005). *The rational imagination: How people create alternatives to reality*. Cambridge, MA: The MIT Press.

Clarke, R. (2014). *Omissions: Agency, metaphysics, and responsibility*. Oxford University Press.

Cummins, D. D., Lubart, T., Alksnis, O., & Rist, R. (1991). Conditional reasoning and causation. *Memory & cognition*, 19.

Gureckis, T. M. et al. (2015). psiTurk: An open-source framework for conducting replicable behavioral experiments online. *Behavior Research Methods*, 1-14.

Goldvarg, E. & Johnson-Laird, P. (2001). Naïve causality: A mental model theory of causal meaning and reasoning. *Cognitive Science*, 25.

Goodwin, G.P., & Johnson-Laird, P.N. (2005). Reasoning about relations. *Psychological Review*, 112.

Henne, P., Pinillos, A., & De Brigard, F. (2016). Cause by omission and norm: Not watering plants. *Australasian Journal of Philosophy*.

Jeffrey, R. (1981). *Formal logic: Its scope and limits* (2nd Ed). New York: McGraw-Hill.

Johnson-Laird, P.N. (2006). *How we reason*. NY: OUP.

Johnson-Laird, P. N., & Byrne, R.M.J. (1991). *Deduction*. Hillsdale, NJ: Erlbaum.

Johnson-Laird P.N., & Byrne R. (2002). Conditionals: A theory of meaning, pragmatics, and inference. *Psychological Review*, 109.

Juhos et al. (2012)

Khemlani, S., Barbey, A., & Johnson-Laird, P. N. (2014). Causal reasoning with mental models. *Frontiers in Human Neuroscience*, 8, 849.

Khemlani, S., & Johnson-Laird, P. N. (2013). The processes of inference. *Argument & Computation*, 4, 1-20.

Khemlani, S., Orenes, I., & Johnson-Laird, P.N. (2012). Negation: a theory of its meaning, representation, and use. *Journal of Cognitive Psychology*, 24.

Nickerson, R. S. (2015). *Conditional reasoning: The unruly syntactics, semantics, thematics, and pragmatics of "If"*. New York Oxford University Press.

Paolacci, G., Chandler, J., & Ipeirotis, P. G. (2010). Running experiments on Amazon Mechanical Turk. *Judgment and Decision Making*, 5.

Peirce, C.S. (1931-1958). *Collected papers of Charles Sanders Peirce*. 8 vols. C. Hartshorne, P. Weiss, and A. Burks, (Eds.). Cambridge, MA: Harvard University Press.

Schaffer, J. (2005). Contrastive Causation. *Philosophical Review*, 114.

Wolff P., Barbey A., Hausknecht M. (2010). For want of a nail: how absences cause events. *Journal of Experimental Psychology: General*, 139.