

## Causation: Empirical Trends and Future Directions

David Rose (Rutgers University) and David Danks (Carnegie Mellon University)

[*Forthcoming* in *Philosophy Compass*]

**Abstract:** Empirical research has recently emerged as a key method for understanding the nature of causation, and our concept of causation. One thread of research aims to test intuitions about the nature of causation in a variety of classic cases. These experiments have principally been used to try to resolve certain debates within analytic philosophy, most notably that between proponents of transference and dependence views of causation. The other major thread of empirical research on our concept of causation has investigated the role that normative considerations play in causal judgments. These experimental results suggest that philosophical accounts of our concept of causation should take a broader view of what might be relevant. For both lines of research, we describe some of the significant experiments and outline key philosophical morals that have been drawn, all while pointing out various limitations. We conclude by considering other kinds of empirical research that should be philosophically interesting for those studying the nature of causation and our concept of causation. In particular, we point towards the need for philosophical research about causal perception, causal reasoning, and causal learning, as well as ways in which this research could play a role in prescriptive metaphysics.

A hot topic among both psychologists and empirically oriented philosophers is the nature of our concept of causation, as well as the ways in which we learn about and use causal relations in the world. One significant thread of research begins with debates within analytic philosophy about the metaphysics of causation. In that literature, there are two dominant views of causation—transference and dependence—which both have substantial support from intuitions about various “test cases.” This evidential stalemate has led some to suggest that we might possess two concepts of causation. One line of empirical work aims to avoid this clash of intuitions among proponents of both transference and dependence views of causation by determining the intuitions and responses of everyday folk. Interestingly, those experiments appear to provide some support for causal pluralism. A second focus of empirical research has been on the role that normative considerations play in causal cognition. Standard philosophical accounts of causation have largely ignored possible influences of normative considerations, such as moral judgments about an actor or beliefs about statistical features of some situation. In contrast, multiple experiments have shown that such considerations do influence people’s judgments, which raises questions about both our folk concept of causation and also the connections between that folk concept and causation as a metaphysical relation “in the world.”

In Section I, we focus on the debate that has primarily occupied analytic philosophers and review empirical work suggesting causal pluralism. Section II turns to research focusing on the role of normative considerations in causal cognition. In Section III, we draw attention to multiple, relatively unexplored avenues of empirical research that are likely to be fruitful in trying to improve our understanding of causation.

### I. The nature of causation

#### *1.1 Causation in analytic philosophy*

At a high level, there are two dominant theories of causation in analytic philosophy: transference and dependence theories of causation. Transference theories of causation are based on the idea that *C* causes *E* just when there is an appropriate *transfer* of some physical quantity between *C* and *E*. The paradigmatic case of causation for a transference theory is the collision of billiard balls: the cue ball hitting the eight ball *causes* the eight ball to move in a particular direction because the cue ball transfers some of its momentum to the eight ball.<sup>1</sup> Different transference theories arise from variations in the relevant quantity, which transferences are “appropriate,” and even what counts as transference,<sup>2</sup> but they all share this common focus on the exchange of some physical quantity.

In contrast, dependence theories of causation center on the idea that *C* causes *E* just when *E* depends on *C* in an appropriate, often counterfactual, manner. A paradigmatic case of “dependence causation” is a neighbor’s failure to water one’s plants (when one is on vacation) causing the plants to die. Clearly, there is no transference between the neighbor and the plants (which is exactly the problem!), but the death of the plants still depends on the neighbor’s (in)action: if she had watered them, then they would have survived. As shown by this example, dependence theories are typically contrastive: they depend on whether *E* would differ if *C* were to be different (hypothetically or counterfactually), and so one must specify the alternatives for *C*. There are many varieties of dependence theories of causation corresponding to the nature of the dependence, which other aspects of the world (if any) should be held constant, and so forth,<sup>3</sup> but they all share this critical core.

The primary strategy in analytic philosophy for deciding between these theories has been to consider one’s intuitive reactions to a range of cases and to favor the theory that is better at predicting and explaining those intuitions. For example, dependence theories can straightforwardly handle omission cases such as one’s plants dying because one’s neighbor failed to water them.<sup>4</sup> In contrast, transference theories cannot capture our intuitions about these cases since there is no transference of any physical quantity, which is precisely what makes them *omission* cases. These cases have been taken by many to be compelling evidence against transference theories as a class (e.g., Beebe, 2004; Godfrey-Smith, 2010; Hall, 2000, 2004; Longworth, 2006; Mellor, 1995).

Proponents of transference theories instead point to preemption cases such as the following:

Billy and Suzy each throw a rock at a bottle. Suzy’s rock arrives first, shattering the bottle, and Billy’s rock arrives just a split second later. Had Suzy’s rock not shattered the bottle, Billy’s still would have.

In this case (and others like it), there is wide agreement that Suzy’s throwing the rock is the actual cause of the bottle shattering, even though it seems that the appropriate dependence fails to hold: the bottle would have shattered even if Suzy had not thrown the rock (since Billy’s rock would have broken it).<sup>5</sup> In general, preemption cases usually lead to the intuition that “the cause” is whatever has the appropriate physical connection with the effect, even if the effect does not exactly depend on that factor. Thus, such cases are taken to show that some transference theory must be correct, rather than any dependence theory (e.g., Hall, 2004; Godfrey-Smith, 2010; Longworth, 2006).

---

<sup>1</sup> Though physical contact is, of course, not necessary for appropriate transference.

<sup>2</sup> See, for example, Dowe (1992, 2000), Fair (1979), Salmon (1984) and Sober (1984).

<sup>3</sup> See, for example, Lewis (1973), Eells (1991), Eells and Sober (1983), Hitchcock (2001), Salmon (1984), Yablo (2002, 2004) and Woodward (2003).

<sup>4</sup> Of course, a dependence theory might also make various counterintuitive claims, such as that the President’s failing to water one’s plants is also an actual cause of them dying.

<sup>5</sup> There are many attempts by dependence theorists to accommodate preemption cases, such as the holding fixed theories of Hitchcock (2001), Yablo (2002, 2004), Woodward (2003), and Lewis’ (1986) discussion in “Postscripts to Causation.” Longworth (2006) presents further cases aimed at undermining these amended dependence theories.

There seems to be agreement that a theory of causation should explain widely held intuitions (else it is not a theory of “causation”), but neither theory-type succeeds in explaining all strong intuitions. One reaction is to argue that both are actually correct: that is, to argue that some version of causal pluralism is the best account of causation (Hall, 2004; Hitchcock, 2003; Godfrey-Smith, 2010). For example, Hall (2004) proposes that there are really two types of causal relations in the world—dependence and transference (what he calls ‘production’)—and that we have two corresponding concepts of causation. These concepts are triggered by salient features of a case, and so we can have seemingly inconsistent intuitions across cases, all while having consistent intuitions within each case.

### 1.2 Empirical tests of intuitions

Given the heavy reliance on intuitive judgments in these areas of analytic philosophy, one natural line of empirical research is to determine whether intuitions in these philosophically interesting cases are widely shared. Lombrozo (2010) explored people’s causal intuitions in cases of double prevention, which often elicit quite strong-but-opposed intuitions among philosophers. Consider the following case of double prevention:

Assassin A aims a gun at President P. Bodyguard B sees A and rushes to protect P. Assassin Z, however, deliberately stops B from reaching P, thereby allowing A a clear shot. P is shot by A and dies as a result.

This is a double prevention case because Z prevents B from preventing A’s action. Clearly, A is a cause of P’s death; the interesting question is whether Z is also a cause of P’s death. Transference theories—and many proponents of those theories—typically answer “no,” since nothing is transferred from Z to P;<sup>6</sup> dependence theories and theorists say “yes,” since P would presumably have survived if Z had acted differently (since B would have been able to protect P). Double prevention cases can thus nicely serve as a litmus test to determine whether one is reasoning using transference or dependence criteria.

Lombrozo reasoned that judgments about double prevention cases might depend on whether the “double preventing” action (e.g., Z’s prevention of B) was understood mechanically or teleologically. To manipulate the participants’ construal, the double preventer was described as acting either accidentally (suggesting a mechanical reading) or deliberately (implying goal-directed action). Goal-directedness implies that the outcome is what matters, not the manner in which it arises, and so dependence is the relevant criterion. In contrast, mechanisms frequently work only one way, and so the focus is on the proper (physical) connections between the components of the mechanism, which prompts a focus on transference.

Lombrozo (2010) used several different versions of double-prevention cases<sup>7</sup> and found that this manipulation was strikingly successful: in cases where both agents (e.g., both assassins A and Z) acted intentionally, they were both judged to be equally strong causes, but when both agents acted accidentally, then causal ratings were significantly higher for the “direct” cause (e.g., assassin A) than the double-preventer (e.g., assassin Z). That is, even though the physical facts about the double preventer’s action are identical in both conditions, intentional behavior prompted dependence-based judgments, while accidental behavior prompted transference-based judgments. As Lombrozo herself points out, this result seems to support the view that we have two different concepts of ‘cause,’ in line with Hall’s (2004) proposal.

Walsh & Sloman (2011) similarly aimed to tease apart dependence and transference intuitions, though using a much larger set of vignettes, many of which will be quite familiar to philosophers (e.g., a slightly

---

<sup>6</sup> Though a transfer theorist could grant that this is plausibly a case of *pseudo-causation* (Dowe, 2001).

<sup>7</sup> Her experiments used relatively innocuous cover stories (e.g., juggling balls hitting radio buttons), rather than the morbid ones favored by philosophers.

modified version of Billy and Suzy throwing rocks at a bottle). Their results similarly point to diversity in our concept of causation, both within- and between-individuals. They find a significant effect of “mechanism”<sup>8</sup> information: a factor is judged to be more causal when people know about a physically continuous process connecting it with the outcome. But importantly, this effect is not complete, as dependence relations mattered for many participants. They also find that people do not seem to use the term “prevent” to mean simply “causes not”: that is, a factor can cause the negation of *X* without preventing *X*, which stands in sharp contrast with the way that term is frequently used in debates in analytic philosophy. It is unclear exactly how this finding should be used by philosophers of causation. Interestingly, many of the vignettes in Walsh & Sloman (2011) produced significant variation in judgments: often, the “minority” judgment would be offered by 30-40% of participants. They did not collect many within-participant judgments, though, so it is difficult to determine the exact cause of this variability. The differences between individuals’ concepts of causation remain an important and intriguing open problem.

### *1.3 Limitations of empirical work on intuitive judgments*

While these sets of experiments are an important first start, neither is without problems. For example, Lombrozo (2010) used intentional vs. non-intentional action as a “cue” to the appropriate construal of the situation, but it is possible that it was actually the intentionality (or not) that mattered, rather than any construal prompted by that description. That is, her experiments are consistent with causal monism if intentionality either (a) is sometimes an important feature of the causal relation in the world (so in our concept of causation), or (b) provides additional information about the causal relation that could influence causal judgments (e.g., about its strength). The potential importance of the intentionality of an action is highlighted by recent work by Rose, Danks, & Machery (2011), who found an interaction between reasons and outcomes: an action was judged to be more of a cause of a bad outcome when it was based on a bad (rather than good) practical reason, and vice versa for good outcomes.<sup>9</sup> Walsh & Sloman (2011) do not have the drawback of a single manipulation, but do use quite complicated vignettes, which presents two potential issues. First, some of the vignettes are sufficiently complex that it is reasonable to wonder whether all participants fully understood them, given that participants are often less-than-fully motivated. Second, the use of vignettes typically presupposes that people’s judgments in these cases match those made when provided with other types of evidence, such as sequences of data. As Danks, Rose, & Machery (under review) show, these need not be the same. But regardless of these issues, the empirical results seem to strongly suggest that the straightforward versions of both transference and dependence theories are wrong about our concept of causation. It remains an open question whether we have multiple concepts of causation or a single, highly complex concept, but in either case, people seem to be sensitive to factors that are not grounded solely in objective physical facts (and laws, counterfactuals, etc.) about a particular situation.

## **II. Normative Considerations and Causal Judgment**

A second line of empirical research has recently emerged that demonstrates that various normative considerations—moral judgments, evaluations of statistical frequency or abnormality, or influences of

---

<sup>8</sup> Their notion of ‘mechanism’ is much more general than that proposed by Machamer, Darden, & Craver (2000).

<sup>9</sup> This is exactly what one would expect if intentionality were a cue to a cause’s strength. Rose, *et al.* (2011) also collected blaming judgments and found that the impact of reason quality could not be explained by differential praise and blame judgments.

other norms<sup>10</sup>—have some influence on causal judgments. That is, our concept of causation, and the intuitive judgments on which standard philosophical theories are based, appears to be sensitive to normative considerations that are typically thought to be irrelevant to causation. These influences have been explained largely in three different ways: normative considerations as a necessary part of causal judgment; causal judgments as judgments of normative responsibility; and normative considerations as an “after-the-judgment” bias in our causal reasoning.

### 2.1 Normative Considerations, Counterfactuals, and Causation

Hitchcock & Knobe (2009) and Knobe (2009) have argued that causal judgments depend in part on counterfactual judgments which themselves depend on normative judgments, and hence normative judgments play a necessary role in causal judgments. That is, this theory holds that there is nothing improper about the influence of normative considerations on causal judgments; rather, such influences reflect the proper method of causal judgment. Multiple theoretical arguments have been provided in support of this view (Hitchcock & Knobe, 2009). In addition, experimental evidence for this argument has come from studies on a wide range of normative considerations, including statistical norm violations (Hitchcock & Knobe, 2009), moral norm violations (Knobe & Fraser, 2008; Hitchcock & Knobe, 2009), and violations of norms of proper functioning (Hitchcock & Knobe, 2009).

To take just one example, consider the Pen Case (Knobe & Fraser, 2008; Hitchcock & Knobe, 2009): a department receptionist is left without any pens (the effect) because an administrative assistant and a professor each took one (the two causes). The effect only occurs if both causes do, and so symmetry considerations suggest that they should be regarded as equally causal in the production of the effect. In contrast, Knobe & Fraser (2008) find that the individual who is not allowed to take pens (i.e., the norm violator) is judged to be significantly more of a cause. This finding is explained on this theory by noting that norm violations make certain counterfactuals salient (e.g., alternatives in which the norm is not violated) and so causal judgments based on such counterfactuals will necessarily reflect this.<sup>11</sup>

### 2.2 Causal Judgments as Responsibility Judgments

The second view of the role of normative considerations in causal judgments holds that causal judgments are not, strictly speaking, about ‘causation’ as analytic philosophers have thought of it. Sytsma, Livengood, & Rose (forthcoming) argue that causal language, at least with respect to the actions of agents, largely reflects concerns with normative responsibility. In other words, they argue that ordinary people’s ascriptions of causation to the actions of some agent typically mean only that the agent is normatively *responsible* for bringing about some outcome.

They used this semantic hypothesis as motivation to reconsider the role that statistical norms play in causal attributions. Knobe & Fraser (2008) showed statistical norm violations (i.e., statistically atypical actions) are judged as more causal than statistically typical actions. Sytsma, *et al.* argued that this claim is ambiguous about whether one means agent-level statistical norms—what is typical for some *particular* agent—or population-level statistical norms—what is typical within some population of which a particular agent is a member. If causal judgments are actually responsibility judgments about a particular individual, then agent-level should influence causal attributions: in particular, agent-level atypical behaviors should be excused as accidents or aberrations for which the agent is less responsible, and so

---

<sup>10</sup> Unless otherwise qualified, we will use ‘normative’ in a broad way to refer to many different types of norms, including moral, statistical, and social ones.

<sup>11</sup> Hitchcock & Knobe (2009) suggest a functional explanation for this behavior: namely, that actions to stop a norm violation are typically the best or most useful interventions in these cases. They recognize, however, that this explanation is quite speculative given the current evidence, though obviously intriguing.

lead to lower causal ratings. In contrast, population-level statistical norms should have relatively little effect, since “what everyone else does” is irrelevant to whether one is responsible for some outcome.

Sytsma, *et al.* used variations on the Pen Case in which there was no mention of pen-taking rules, but in which different people have different agent-level and population-level typical actions. They found that population-level statistical norms had no effect on people’s causal judgments but that agent-level statistical norms did have an effect: people judged that the individual whose behavior was agent-level typical was judged to be much more of a cause of the problem than the individual who behaved agent-level atypically. Moreover, these results speak against the hypothesis of Knobe & Fraser (2008) and Hitchcock & Knobe (2009), as they provide a case in which a statistical norm violation actually led to a *decrease* in causal attribution.

### 2.3 Normative Considerations as a Bias in Causal Judgments

Both of the preceding accounts propose that the influence of normative considerations is appropriate, though for different reasons; that is, both hold that normative considerations *should* influence causal judgments. In contrast, Alicke, Rose, & Bloom (forthcoming) argue that the influence of normative considerations on causal judgments is simply a bias (see also Alicke, 1992). Their Culpable Control Model first proposes that negative outcomes can trigger a spontaneous negative evaluative reaction in an observer, whether because of central (e.g., whether the agent desired or foresaw that her action would bring about the negative outcome) or peripheral (e.g., the social status or race of the agent) factors. This negative evaluation then prompts a blaming response in the observer, who adjusts her causal judgment in order to justify that blaming reaction (since blame for outcome *O* is only justified when the agent actually caused *O*). As a result, negative outcomes lead to retroactive revision of causal judgments in order to minimize the cognitive dissonance that arises when one blames an individual *A*, even though *A* is judged to be a minimal cause, or not a cause at all.

Several lines of empirical support have been provided in favor of this view. Alicke (1992) found that, when all physical facts were held constant, an agent who brought about a negative outcome for a socially undesirable motive was both blamed more and judged to have a greater causal role than one who acted on a socially desirable motive. Alicke, *et al.* (forthcoming) found the same pattern of findings in a much wider range of cases. In one such study, they presented a vignette about a student who either cheated or did not, and either by himself or in a group. In all cases, the teacher graded on a curve with a fixed number of A’s, and the cheating student got the lowest A. As a result, another student (who did not cheat) got a B, and was rejected from medical school as a result. Alicke *et al.* found that cause and blame ratings only depended significantly on whether the student acted alone or as part of a group when he did *not* cheat. Specifically, the student received significantly less blame and causation only when he chose not to cheat on his own. And in another, related study, Alicke *et al.*, found that the impact of norm violations on causal judgments was mediated by judgments of blame, which suggests that blaming is potentially a cause of those very causal judgments. These findings put pressure on the theories of Sytsma, *et al.* and Hitchcock & Knobe as they suggest that causal judgments are subject to irrelevant<sup>12</sup> influences.

### 2.4. Limitations on Normative Consideration Experiments

These experimental results about causal judgment have been used as evidence for the broad claim that the concept of causation is fundamentally tied up with normative considerations, or that the influence of such considerations is ubiquitous in causal cognition (e.g., Hitchcock & Knobe, 2009; Knobe, 2009; Alicke, 1992). Recent work by Danks, Rose, & Machery (under review) has criticized this view, as the extant work has primarily come from studies using vignettes to prompt conscious judgments. Causal cognition is not composed solely of classificatory causal judgments, nor is it based solely on textual vignettes. Rather,

---

<sup>12</sup> At least, irrelevant to the task of determining whether one factor caused another.

causal cognition encompasses many different processes—causal learning, causal inference, causal perception, and causal reasoning—involving representations generated from a variety of information and data.

Danks, *et al.* investigated whether normative considerations play a role for some of these other types of causal cognition. They provided participants with two sequences of observational data, each about a potential cause and a target effect. The control scenario was morally neutral, while the moralized scenario focused on an individual who was trying to destroy a cure for cancer. After seeing a sequence of cases (i.e., whether the potential cause or target effect occurred in each of a set of individuals), participants reported the causal strength of the potential cause, as well as the blameworthiness and knowledge state of the morally bad actor for the moralized scenario. Danks, *et al.* varied the sequence statistics between participants, and found (in line with numerous other psychological experiments) that changes in case frequencies produced significant differences in causal strength ratings. More interestingly, they found *no* effect of normative considerations on those ratings: causal inference from a sequence of cases does not seem to be subject to the same influences of normative considerations as causal judgment from vignettes. The exact scope of these influences on causal cognition thus remains an open question, but clearly normative considerations do not influence all causal cognition equally.

### III. Future Directions and Unexplored Avenues

#### 3.1. New Empirical Avenues

As noted in the previous section, most of the extant empirical philosophical work on causation has focused on causal judgments based on vignettes. While this research is undoubtedly philosophically relevant, one could look to other aspects of causal cognition to learn more about the philosophically relevant aspects of this concept. Much of the psychological research on causal cognition has done exactly that, focusing on such aspects as causal perception, causal reasoning, and causal language.

Causal perception refers to the relatively automatic, direct perception of one event causing another (see Rips, 2011 or Scholl & Tremoulet, 2000 for a review). The canonical example is the so-called “launching effect” (Michotte, 1963) in which one rectangle moves towards another, and the second begins moving when there is contact from the first; perceptually, the first rectangle “hits” the second and makes it move. Causal perception can also arise for “social causation”—e.g., one individual “chasing” another (Heider & Simmel, 1944)—and more complex physical phenomena—e.g., explosions triggered without contact (White & Milne, 1997). At the same time, causal perception is both highly context-dependent (e.g., Choi & Scholl, 2004; Scholl & Nakayama, 2002) and connected with judgments about the intentional states of possible agents (e.g., Leslie, 1982). Further work is required on the connections between causal perception and our concept of causation (though see, e.g., Danks, 2009; Woodward, 2011).

Causal knowledge and judgments are also used to reason about systems in the world, such as determining which factors might be causally relevant, or which actions will likely to lead to a desirable outcome. Research on causal reasoning has largely focused on whether causal knowledge is represented as something like a causal Bayesian network (e.g., Gopnik, *et al.*, 2004) and whether that knowledge is simply correlational or also includes information about mechanisms (e.g., Ahn, Kalish, Medin, & Gelman, 1995). Although there has been some use of this research in philosophical circles, the connections are quite complex. For example, the notion of ‘mechanism’ that is used by Ahn, *et al.* (1995) is quite different from the philosophical notion of ‘mechanism’ found in, e.g., Machamer, *et al.* (2000). And the fact that people can use causal knowledge to design appropriate interventions (e.g., as children do in Gopnik, *et al.*’s experiments) does not thereby show that our concept of causation is interventionist (in the sense of Woodward, 2003) in any deep or necessary way.

Finally, one could look directly at the ways that people use causal language in everyday life. This line of research has found that many instances of causal language can be best understood as carrying information about “force dynamics,” either actual physical forces or metaphorical social or psychological “forces” (Talmy, 1988; Wolff, 2007; Wolff & Song, 2003). Our use of terms like ‘hinder’, ‘assist’, ‘enable’, and so forth all seem to carry information about default states and the behaviors of components of the causal system under consideration. These results seem to support the role of normative considerations, since these default states and behaviors are captured in (at least) statistical norms. However, there are numerous questions about the generalizability of these results to situations in which causes cannot be easily understood as having “forces” or “strengths,” whether physical or metaphorical. Thus, serious philosophical work remains to be done before these results can be used to better explicate our concept(s) of causation.

### *3.2. Empirical Work and the Metaphysics of Causation*

Empirical work on causation has, unfortunately, paid little attention to how that work might bear on the metaphysics of causation. Both Goldman (1987, 1992, 2007) and Paul (2010) have suggested various ways that cognitive science might be relevant to metaphysical theorizing. Following Goldman (1992) we can usefully distinguish between *descriptive* and *prescriptive* metaphysics. Descriptive metaphysics is “the discipline that seeks to describe and understand...folk ontology...seek[ing] to lay bare its content and understand its roots and sources” (Goldman, 1992, p. 35). Prescriptive metaphysics, by contrast, “would try to tell us what metaphysical commitments we *ought* to adopt, given the best available science and philosophy” (p. 35). The empirical work we have focused on falls within the domain of descriptive metaphysics since that work has tried to understand the nature of our everyday causal cognition. At the same time, this descriptive work is clearly relevant even for “standard” metaphysics, as it helps to specify and articulate the intuitive judgments that any theory of causation (in the world) must explain.

A key issue for future research on causation and cognitive science will be finding ways for that empirical research to play an important role in prescriptive metaphysics. Goldman (1992) argued that this would require better “understanding our initial view of the world and our own constitution” (p. 36) so that we can determine and evaluate various possible causal understandings of the world. There are a range of revisionary moves that are potentially open depending on progress on the scientific understanding of our causal cognition, including eliminativism (perhaps we should do without causal language), reductionism (perhaps causal language should be replaced with various sub-types), and error theories (perhaps causal language does not track any actual causal relation). We obviously cannot undertake the enormous task of cognitive science-based prescriptive metaphysics for causation in this paper. But we note in closing that this area of inquiry has been almost completely neglected in empirical philosophical research on our causal cognition, and our concept(s) of causation.



## References

- Ahn, W.-K., Kalish, C. W., Medin, D. L., & Gelman, S. A. (1995). The role of covariation versus mechanism information in causal attribution. *Cognition*, *54*, 299-352.
- Alicke, M. (1992). "Culpable Causation. *Journal of Personality and Social Psychology*, *36*: 368–378.
- Alicke, M., Rose, D., & Bloom, D. (forthcoming). Causation, Norm Violation and Culpable Control. *Journal of Philosophy*.
- Beebe, H. (2004). Causing and Nothingness. In L. A. Paul, E. J. Hall & J. Collins (eds.), *Causation and Counterfactuals*.
- Choi, H., & Scholl, B. J. (2004). Effects of grouping and attention on the perception of causality. *Perception & Psychophysics*, *66*, 926-942.
- Danks, D. (2009). The psychology of causal perception and reasoning. In H. Beebe, C. Hitchcock, & P. Menzies (Eds.), *Oxford handbook of causation* (pp. 447-470). Oxford: Oxford University Press.
- Danks, D., Rose, D., & Machery, E. (under review). Demoralizing causation. Submitted to *Philosophy and Phenomenological Research*.
- Dowe, P. (1992). "Wesley Salmon's Process Theory of Causality and the Conserved Quantity Theory," *Philosophy of Science*, *59*, pp. 195-216.
- Dowe, P. (2000). *Physical Causation*, New York: Cambridge University Press.
- Dowe, P. (2001). A counterfactual theory of prevention and 'causation' by omission. *Australasian Journal of Philosophy*, *79*, 216-26.
- Fair, D. (1979). "Causation and the Flow of Energy," *Erkenntnis*, *14*, pp. 219-250.
- Godfrey-Smith, P. (2010). "Causal Pluralism," in: Hitchcock, C., Beebe, H. and Menzies, P. (ed s.), *Oxford Handbook of Causation*.
- Goldman, A. (1987). Cognitive Science and Metaphysics. *Journal of Philosophy*, *84*, pp. 537-544.
- Goldman, A. (1992). *Liaisons: Philosophy Meets the Cognitive and Social Sciences*. MIT Press.
- Goldman, A. (2007). A Program for "Naturalizing" Metaphysics, with Application to the Ontology of Events. *The Monist*, *90*, pp. 457-479.
- Gopnik, A., Glymour, C., Sobel, D. M., Schulz, L. E., Kushnir, T., & Danks, D. (2004). A theory of causal learning in children: Causal maps and Bayes Nets. *Psychological Review*, *111*, 3-32.
- Hall, N. (2000). "Causation and the Price of Transitivity," *Journal of Philosophy*, *97*, pp. 198-222.

- Hall, N. (2004). "Two Concepts of Causation," in: Collins, J., Hall, N., and Paul, L.A. (eds.), *Causation and Counterfactuals* (pp. 225-276), Cambridge: MIT Press.
- Heider, F., & Simmel, M.-A. (1944). An experimental study of apparent behavior. *American Journal of Psychology*, 57, 243-249.
- Hitchcock, C. (2001). "The Intransitivity of Causation Revealed in Equations and Graphs," *Journal of Philosophy*, 98, pp. 273-299.
- Hitchcock, C. (2003). "Of Humean Bondage," *British Journal for the Philosophy of Science*, 54, pp. 1-25.
- Hitchcock, C., & Knobe, J. (2009). Cause and Norm. *Journal of Philosophy*, 106 (11), 587-612.
- Knobe, J., & Fraser, B. (2008). "Causal judgments and moral judgment: Two experiments." In W. Sinnott-Armstrong (Ed.), *Moral Psychology, Volume 2: The Cognitive Science of Morality*, pp. 441-447. Cambridge: MIT Press.
- Knobe, J. (2009). Folk Judgments of Causation. *Studies in the History and Philosophy of Science* 40:2, 238-242.
- Leslie, A. M. (1982). The perception of causality in infants. *Perception*, 11, 173-186.
- Lewis, D. (1973). "Causation," *Journal of Philosophy*, 70, pp. 556-567.
- Lewis, D. (1986). *Philosophical Papers*, Volume II. Oxford: Oxford University Press.
- Lombrozo, T. (2010). Causal-explanatory pluralism: how intentions, functions, and mechanisms influence causal ascriptions. *Cognitive Psychology*, 61, 303-332.
- Longworth, F. (2006). Causation, Pluralism and Responsibility. *Philosophica*, 77, 45-68.
- Machamer, P., Darden, L., & Craver, C. F. (2000). Thinking about mechanisms. *Philosophy of Science*, 67, 1-25.
- Mellor, D. H. (1995). *The facts of causation*. London and New York: Routledge.
- Michotte, A. (1963). *The perception of causality*. London: Methuen.
- Paul, L. (2010). A New Role for Experimental Work in Metaphysics. *Review of Philosophy and Psychology*, 1, pp. 461-476.
- Rips, L. J. (2011). Causation from perception. *Perspectives on Psychological Science*, 6, 77-97.
- Rose, D., Danks, D., & Machery, E. (2011). The Reasons Model of Actual Causation. Unpublished Manuscript: Rutgers University.
- Salmon, W.C. (1984). *Scientific Explanation and the Causal Structure of the World*, Princeton: Princeton University Press.

- Scholl, B. J., & Nakayama, K. (2002). Causal capture: Contextual effects on the perception of collision events. *Psychological Science, 13*, 493-498.
- Scholl, B. J., & Tremoulet, P. D. (2000). Perceptual causality and animacy. *Trends in Cognitive Sciences, 4*, 299-309.
- Sober E. (1984). "Two Concepts of Cause," *PSA 1984*, vol. 2, pp. 405-424.
- Sytsma, J., Livengood, J., & Rose, D. (forthcoming). Two Types of Typicality: Rethinking the Role of Statistical Typicality in Ordinary Causal Attributions. *Studies in the History and Philosophy of Science*.
- Talmy, L. (1988). Force dynamics in language and cognition. *Cognitive Science, 12*, 49-100.
- Walsh, C. R., & Sloman, S. A. (2011). The meaning of cause and prevent: The role of causal mechanism. *Mind & Language, 26*, 21-52.
- White, PA, & Milne, A. (1997). Phenomenal causality: impressions of pulling in the visual perception of objects in motion. *American Journal of Psychology, 110*, 573-602.
- Wolff, P. (2007). Representing causation. *Journal of Experimental Psychology: General, 136*, 82-111.
- Wolff, P., & Song, G. (2003). Models of causation and the semantics of causal verbs. *Cognitive Psychology, 47*, 276-332.
- Woodward, J. (2003). *Making Things Happen: A Theory of Causal Explanation*, Oxford: Oxford University Press.
- Woodward, J. (2011). Causal perception and causal cognition. In J. Roessler, H. Lerman, & N. Eilan (Eds.), *Perception, Causation, and Objectivity*. Oxford: Oxford University Press.
- Yablo, S. (2002). "De Facto Dependence," *Journal of Philosophy, 99*, pp. 130-148.
- Yablo, S. (2004). "Advertisement for a Sketch of an Outline of a Prototheory of Causation," in: Collins, J., Hall, N., and Paul, L.A. (ed s.), *Causation and Counterfactuals* (pp. 119-137), Cambridge: MIT Press.