Norms in Actual Causation

Jenn McDonald

Forthcoming in Erkenntnis

Abstract Experiments in psychology and experimental philosophy suggest that judgments about actual causation are partially governed by norms: norm violations are more likely to be singled out as causes, while structurally analogous factors that obey the norms are unlikely to be singled out. The norm-sensitivity of causal judgment has, in turn, lent support to a normative analysis of causation itself. In this paper, I question whether the support stands. I articulate and examine two principal reasons support might be so derived. For each, I argue that, in fact, a non-normative analysis is better supported.

§1 Causal Judgment and Metaphysics

A number of experiments in psychology and experimental philosophy, designed to elicit causal judgments of various kinds (Hitchcock & Knobe, 2009; Knobe & Fraser, 2008; Rose et al., 2021), indicate that norms play some role in how we assign causal status to different factors in a situation.[1] In particular, norm-violating factors seem to be substantially more likely to be picked out as having a special causal status – as being 'the' cause of some effect. For example, Knobe and Fraser ran a study that presented the following vignette.

> **Pen Vignette** The receptionist in the philosophy department keeps her desk stocked with pens. The administrative assistants are allowed to take the pens, but faculty members are supposed to buy their own. [A]dministrative assistants typically do take pens. Unfortunately, so do faculty members. The receptionist has repeatedly emailed them reminders that only administrative assistants are allowed to take the pens. On Monday morning, one of the administrative assistants encounters Professor Smith

[1] The term 'factor' is used á là Eells (1988, 1991), but see also (Hitchcock, 2001, p. 362; Menzies, 2004b). It refers neutrally to candidate causal relata, possibly ranging over various kinds of entities (events, states-of-affairs, property instances, facts, propositions, etc.). The choice of relata is inconsequential to the argument.

> walking past the receptionist's desk. Both take pens. Later that day, the receptionist needs to take an important message... but she has a problem. There are no pens left on her desk. (Knobe & Fraser, 2008, p. 441)

Respondents were then asked to report their agreement (or disagreement) with each of the following two claims.

- "'Professor Smith caused the problem.'
- 'The administrative assistant caused the problem.'" (Knobe & Fraser, 2008, p. 441)

Overwhelmingly, respondents agreed that the professor borrowing a pen caused the problem, but not the administrator's doing so. Many other studies have since been run using different vignettes involving a variety of norms.[2] Crucially, the kinds of norms found to induce this effect include highly pragmatic and subjective ones, such as moral, social, or conventional norms, as well as more objective ones, such as statistical norms or norms of proper functioning.

The open question in the psychology literature is what the role of norms is, exactly, in actual causal judgments. One view, dubbed the *normative concept view*, holds that our concept of actual causation is inherently normative – with normative considerations playing a constitutive role in the formation of any causal judgment (Hart & Honoré, 1985; Hitchcock & Knobe, 2009; Knobe, 2009; Sytsma & Livengood, 2021).[3] As an account of our concept of actual causation, this is not so implausible. It would be unsurprising to learn, perhaps, that our ordinary concepts in many areas are inherently shaped by the pragmatic considerations that govern their use and role in our lives. But the normative concept view has been taken to lend support to an increasingly popular view about metaphysics – that a correct analysis of

[2] This includes many examples that avoid using normatively loaded language like 'the problem' in the prompt.

[3] The two other principal views explain the results in ways that need not invoke an inherently normative causal concept. The *norm-sensitive cognitive process view*, for example, holds that our concept is descriptive, yet norms may influence which of several possible causal judgments we endorse. For example, norms may inform which counterfactual alternatives are considered when forming a causal judgment. See (Willemsen & Kirfel, 2019) for an overview.

actual causation must incorporate a normative parameter (Gallow, 2021; Hall, 2007; Halpern, 2008; Halpern & Hitchcock, 2010, 2015; Menzies, 2004a, 2017). The idea here is that whether or not some factor violates a norm is an essential determiner of whether that factor is an actual cause or not of some effect. Not simply whether we would pick it out as such, but whether it is a cause at all. Call this a *normative analysis* of actual causation.

If correct, this is surprising.[4] By incorporating norms into our analysis, we give up on the idea that the relation of actual causation is *real* – that whether the relation holds between two relata is determined wholly independently of how we think or talk.[5] The basic issue is that many of the relevant norms are anthropocentric (social, conventional, moral, etc.) and arguably could not all be given a realist construal. To compound matters, there are regularly competing sets of norms – 'competing' in the sense that the same factor may violate one set of norms while not violating another, where both plausibly hold of the situation (Blanchard & Schaffer, 2017). A given event might be statistically normal but morally abnormal, for example, or statistically normal relative to one reference class but not relative to another. And yet, said event may be a cause relative to one set of norms but not the other(s). So, what counts as a cause will further depend on *which* norms of various kinds are causally relevant. There are two possible responses, and neither recovers realism. One could require a unique determination of relevant norms. However, establishing a unique determination will invariably call upon pragmatic considerations regarding our interests and purposes, which further undermines realism (Blanchard & Schaffer, 2017). Alternatively, causation itself may hold relative to a set of norms. However, since some of these norms are anthropocentric, actual causation remains not entirely mind and language independent.

I believe the step from recognizing the role of norms in causal judgment to incorporating them into our metaphysics has been far too quick. Even granting the normative concept view

[4] Whether it is correct is another question. See especially (Wysocki, forthcoming).

[5] Note that the question of realism about a relation is largely distinct from that of its relata. The mind and language independence of actual causation is consistent with the dependence of some range of events or property instances.

as the best interpretation of the experimental data, there has been remarkably little direct engagement with the question of how this aspect of causal judgment should bear on our metaphysical analysis of causation. I take there to be two principal reasons to infer from the normative nature of causal judgment to the normative sensitivity of causation itself. Both reasons have to do with the need to explain why, for whatever relation is proposed as an analysis, that very particular relation is so qualified to go under the name 'actual causation'. In what follows, I consider these reasons in turn and offer a response on behalf of a non-normative analysis of actual causation.

§2 Paying What Metaphysics Owes to Intuition

The first reason to infer to a normative analysis of causation from a recognition of norm-sensitivity in causal judgment takes into account the way in which an analysis, in general, is defeasibly answerable to pre-theoretic intuition – to causal judgment, in this case. A common test of the viability of an analysis of some concept or phenomenon, ϕ, checks the results of the analysis against ϕ-intuitions in a variety of hypothetical cases. Widely known as the *method of cases*, an analysis is good to the extent that it accords with intuition across a range of cases, while a violation defeasibly counts against it (Pust, 2019). Call the underlying methodological principle the *intuition desideratum*. This is a standard desideratum in contemporary analyses of causation. For example, Paul and Hall write,

> [A]nalyses that aim to develop a more theoretically useful notion of 'cause' (useful, again, either to metaphysicians or to scientists) should take causal intuitions as defeasible guides to potentially interesting and important features of our causal concept or a causal relation… (2013, p. 2)

Insofar, then, as causal judgment is sensitive to norms, there is defeasible pressure on an analysis of causation to also be so sensitive. I take this to be the first principal reason for adopting a normative view of causation – that it satisfies the intuition desideratum.

However, this is not the only way intuition puts pressure on a metaphysical analysis. Another compelling desideratum of an analysis of ϕ is that that which is identified with ϕ have the properties we intuitively take ϕ to have. Call this the *property desideratum*. Intuitively, causation is something independent of minds and languages; a relation that would hold of real things in the world in roughly the same way whether humans had evolved or not. We've seen, however, that the relation described by a normative view doesn't have this property, and so the view violates this desideratum.

An alternative view that satisfies the intuition desideratum while retaining realism about causation would therefore be attractive.[6] One promising alternative is a kind of contrastivism called *Causal Relativism* (McDonald, 2022). According to Causal Relativism, actual causation holds relative to a background space of possibilities, which gives the specific variety of ways situational factors – including putative causal factors – could have gone differently. If a space of possibilities is modeled by a set of possible worlds, then the view claims that the particular relation that holds between two factors, *c* and *e*, will be different against one set of possible worlds than against another.[7] Further, there is no coherent way to model something like a universal space of possibilities – this, due to constraints certain factors being the case place on whether other factors can hold. For example, if it's not possible to have a second train travelling down the tracks, then it's not possible for two different sets of train tracks to be occupied at the same time. Given all this, the full story of whether causation holds between *c* and *e* is determined by the relation holding between them against some specification of background possibilities.

This underlying view of causation can then be exploited to satisfy the intuition desideratum without giving up on realism, and so also satisfy the property desideratum. First, the view

[6] Perhaps a case could be made that intuitions about whether a particular thing causes something else are more reliable (in general?) than those about what holds generally of causal relations. If so, the intuition desideratum would be more important than the property desideratum. My point is simply that a view which could satisfy both is preferable.

[7] The exact nature of this relation doesn't matter for the purposes of the argument. One could treat it as a kind of complex counterfactual dependence (reductive or non-reductive) or as a primitive type-level causation.

postulates that a specification of a space of possibilities is an implicit component of causal claims. That causal claims are context-sensitive in some way is highly plausible (Dretske, 1977; Schaffer, 2012). The proposal is that judgments of actual causation operate in part by implicitly selecting a background possibility space, relative to which the causal judgment holds. Call a *normal possibility space* one in which actual factors that violate norms are replaced by factors that instead cooperate, while any actual factors that obey the norms are held fixed (all relative, of course, to some coherent set of norms). Then, for a normative concept view, our concept of causation invokes only normal possibility spaces. For a norm-sensitive cognitive process view, normal possibility spaces guide our selection of counterfactual alternatives. Either way, we judge that an actual factor is a cause if and only if it makes a difference to the putative effect relative to a (particular/privileged) normal possibility space.

To see how this works, consider again the Pen Vignette. The professor's borrowing a pen violates a norm – the social/moral norm prohibiting professors from taking pens. The salient space of possibilities then will be one which corrects the violation – namely, one in which the professor did not borrow a pen. Alternative possibilities for factors that are already in cooperation with norms are not considered. Since the administrator's borrowing the pen obeys the stated norms, an alternative action on the administrator's part is not considered. Did the professor cause the problem? The normal possibility space is one in which the professor did not take a pen but the administrator still did. In such a case, there is no problem. So, the professor is a cause.

How about whether the administrator caused the problem? There are two ways our judgment that they didn't could be explained. First, we consider the normal possibility space; the one in which the professor did not take a pen but the administrator still did. Since, relative to this alternative, the administrator could not have done differently, their action is obviously not a cause. One might insist, however, that causal judgments always involve considering the (perhaps improbable) possibility of the putative cause being different. Then, comparing what actually occurs simply to a normal possibility space cannot exhaust what goes on in causal judgment. A second way, then, requires that when the putative cause is

normal there be some other consideration or cognitive process that comes into play. In line with Menzies (2004a, 2004b), but constraining it to our psychology, I would suggest that the relevant comparison in such cases becomes that between what occurs in a normal possibility space and what would occur in this normal space had the putative cause been different.[8] Per the Pen Vignette, alter the normal possibility space only so that the administrator does not take a pen. Given this comparison, whether the administrator takes a pen (in the normal possibility space, where neither does the professor) or does not take a pen (in the slightly altered space, where again neither does the professor), there is no problem. Thus, the administrator is not a cause of the problem.

This style of explanation of causal judgment – one that involves our selecting or paying special attention to norms – should be familiar from the literature. What causal relativism contributes is not a novel explanation of how norms affect causal judgment, but an underlying realist metaphysics compatible with this kind of explanation.[9] It allows our causal judgments to endorse causal relations which are implicitly relativized to spaces of possibilities that reflect the norms in place. But this selection process being governed by norms is compatible with a realist construal of the underlying relativized causal relations. This is because our causal judgments don't tell the full story. Causal relations that hold relative to normal possibility spaces are just some of the total network of causal relations in the world. Lots of other causal relations hold relative to other, non-normal, possibility spaces. For example, the administrator taking a pen *is* a cause of the problem relative to the possibility space in which the administrator doesn't take a pen but the professor still takes one. Causal relativism denies that any of the many relativized causal relations, any of the possibility spaces relative to which causal relations hold, is metaphysically privileged over any others. But it allows that we pick some out and focus on them due to their being of

[8] Yablo (1992, p. 274) makes a similar suggestion in another context.

[9] As a result, causal relativism can be treated as supplementary to various proposals which are either quiet about the underlying metaphysics or otherwise leave it open. See, for example, (Blanchard & Schaffer, 2017; Halpern & Hitchcock, 2015; Icard et al., 2017). The view also goes well with any of the three principal accounts of how norms influence causal judgment, as laid out in (Willemsen & Kirfel, 2019).

interest or of use. That these ones are of use of course depends on us. But that they hold in the world is independent.

Thus, when considering what a metaphysical analysis owes to intuition, a non-normative analysis of causation – causal relativism, in particular – does better justice to how we intuitively conceive of actual causation than does a normative one.

§3 Explaining Our Concept of Actual Causation

The second reason to infer from the normativity of causal judgment to that of causation is the need to explain why, for whatever analysis of actual causation one proposes, we should have come to have a concept corresponding to that relation. A normative analysis has a nice explanation for this: there is simply no benefit to our having a notion like actual causation *unless* it is inherently normative (Hitchcock & Knobe, 2009).

To see this, first note that knowledge of causal structure is essential to the provision of satisfying explanations and accurate predictions. However, many think that knowledge of *type-level* or *general* causation – the kind involved in claims of the form "Smoking causes cancer" – is sufficient to do this work so long as it's supplemented with enough information about the concrete situation. So, what work is left for a concept of *actual* or *singular* causation – the kind involved in claims of the form "Jim's smoking caused his cancer" – to do? Another way to put this is in terms of *token causal structure* (Hitchcock, 2017; Hitchcock & Knobe, 2009). The token causal structure of a particular situation specifies a network of dependency relations between various situational factors without specifying what actually causes what. Such structure is well captured by structural equation models, although reference to models isn't essential.[10] But actual causation judgments are something over and above an

[10] Structural equation models represent factors using values of variables (with alternative ways those factors could have gone filling in the remaining values of the variable). Dependencies between factors are represented using asymmetric functional equations defined over the variables. These equations say precisely how the values of different variables determine the values of others. See, among others, (Pearl, 2000/2009).

appreciation for token causal structure. This can be seen in cases of preemption. Say we have a model that accurately represents a situation in which Suzy throws a rock at a window and it shatters, Billy throws a rock shortly after Suzy, and Billy's rock passes through the space where the window used to be moments after it is shattered by Suzy's rock. We might think the model tells us everything we need to know about what occurs and what would have occurred had things been different in particular ways. In addition to this, though, we judge Suzy's throw to be the actual cause of the window shattering, where Billy's throw is not an actual cause. It is notoriously difficult to identify the precise role that Suzy plays that makes only her throw the actual cause, consistent with actual causation judgments in other cases.

One possible way forward asks after the purpose served by an appreciation of actual causation over and above that of token causal structure. Say we already have an understanding of the token causal structure of a situation. The only possible additional benefit, or so argues Hitchcock (2017), is that actual causation captures information about path-specific effects in the context of goal-oriented inquiry. This, in turn, explains why norms are involved – in striving towards our goals we look for interventions available to us *within* the confines of the various normative constraints on our action. If this is correct, though, then the proponent of a non-normative analysis of actual causation must concede that we do not have a concept corresponding to their posited relation. And if not, then the non-normative analysis has a significant, additional burden to provide good grounds for why the non-normative relation should be called 'actual causation' – grounds comparative to those provided by a normative analysis via the above argument.

But the question is loaded. It seems to presuppose that any explanation of our having a concept would be *functional* – providing a good story about what purpose that concept serves. Perhaps the reasoning behind this lies on analogy with natural selection – that a concept would only have survived were it to accrue benefit to the conceiver. However, as a biological and psychological phenomenon, concept acquisition is a complex enterprise. In fact, the psychology literature provides an easy, alternative, *etiological* explanation: the concept of actual causation derives from a fundamental feature of causal learning – reliance on singular cues to causation. In what remains, I first illustrate and explain the nature of

singular cues and of their role in causal learning. I then lay out a developmental proposal about our concept of actual causation that, if correct, better supports a non-normative analysis than a normative one – whether or not our mature concept of actual causation is inherently normative.

§4 Reliance on Singular Cues

There is ample evidence suggesting that a key component of causal learning lies in the recognition of particular, singular cues. Singular cues are those that are fully present in a single, particular situation and whose uptake requires only observation of that singular situation. These include spatiotemporal contiguity and temporal order (Bullock et al., 1982; Burns & McCormack, 2009; Lagnado & Sloman, 2004, 2006; Leslie, 1982, 1984; McCormack et al., 2015; Mendelson & Shulz, 1976; Michotte, 1946; Oakes, 1994; Rottman et al., 2014; Shultz, 1982; White, 1988; Young & Sutherland, 2009), and, arguably, some kinds of observations of mechanistic or generative relations (Ahn et al., 1995; Campbell, 2020; Michotte, 1946; Shultz, 1982; Shultz, Altmann, et al., 1986; Shultz, Fisher, et al., 1986; White, 1988, 1995; Wolff, 2007). Thus, singular cues encompass cues relevant to causal perception as well as some more relevant to causal inference.[11]

To illustrate the role of one such cue – spatiotemporal contiguity – consider experiments on the launching effect, as first discovered by Michotte (1946). The phenomenon of the launching effect refers to our experience of phenomenal appearances as distinctively causal. In the prototypical experiment, one object (usually presented on a screen) approaches a stationary second object. Upon contact, the first object stops moving and the second object begins moving along the same trajectory as the first. This interaction is universally experienced as causal. More precisely, the first object is seen as 'launching' the second.

[11] I assume a distinction between these two areas as outlined by Danks: "Causal perception consists of the relatively automatic, relatively irresistible perception of certain sequences of events as involving causation. Causal inference … consists of higher-level causal learning that is based largely on statistical relationships." (2009, p. 448)

Relevant for my purposes is the way in which disruptions of the spatiotemporal contiguity of the interaction – either by introducing a spatial gap (so the objects don't in fact touch), or a temporal delay (so the second object is stationary for some period after the first object makes contact), or both – disrupt our experience of the interaction as causal. Disruption of spatial contiguity especially, but sufficiently large disruptions of either type leads to our no longer experiencing the interaction as causal. This is evidence that spatiotemporal contiguity is a cue for causal learning.[12]

Crucially, recognition of singular cues does not require consideration of other situations, such as comparison to situations with a similar profile. But normative considerations do. These only apply at the type-level. Statistical norms report type-level facts (such as patterns in the data, averages, etc.), injunctive norms report constraints on general populations or on behavior of a certain type, while norms of proper functioning describe the role of something, defined relative to a type of mechanism, organism, or artifact. So, while the application of a normative concept requires typing a situation in some way, the uptake of singular cues does not require anything beyond the particular situation. Thus, causal learning via singular cues can run independently of normative considerations.[13] This allows room to argue for a non-normative causal concept whose content is a singular relation holding between two particular things.

To advance this argument, I propose that a concept of actual causation as a non-normative, singular relation – associated with our reliance on singular cues – is ontogenetically primary.[14] From here, the concept develops over time and in consideration and

[12] Note that I am not taking a stand on the debate around the *directness* of causal perception, in which the launching effect prominently figures. See (Rips, 2011) for good discussion of that debate.

[13] That this process may be informed by normative considerations at some stage – perhaps as a way of selecting salient singular cues – is beside the point here. All that matters dialectically is that uptake via singular cues can be independent of normative considerations.

[14] I borrow the expression, 'ontogenetically primary', from Carey (2009). Note that by 'primary,' I mean a concept associated with reliance on singular cues is *developmentally prior*, rather than that it conceptually grounds our mature concept. I remain neutral about how our mature concept is structured (except when

accommodation of a wider range of data, such as counterfactual dependencies and statistical or normative considerations. Eventually, it evolves into our mature, adult concept of actual causation.[15]

Support for this proposal derives from the fact that reliance on singular cues holds a particular kind of privilege in causal learning. Principally, they are relied upon in earliest causal learning (Mascalzoni et al., 2013).[16] This implies that particular observations of singular situations suffice for causal learning. Some evidence also suggests that we have a preference for acquiring more information about the single situation under causal inquiry, rather than information about other similar situations (White, 1989, 2014).

But how does this support a non-normative analysis? Suppose first that our mature concept of actual causation is inherently normative. The proponent of a normative analysis may then be tempted to argue that what matters when considering implications for metaphysics is the content of our mature concept, not the content of various developmental stages of said concept. Thus, if our mature concept is normative, as granted and as the normative concept view holds, then surely an inherently normative causal relation has better claim to being called 'actual causation' than a non-normative one. That our proto-concept is non-normative is beside the point, this argument might run.

This response helps to distinguish between two meta-metaphysical questions. (1) What constitutes good grounds for positing the existence of a kind of relation in the world? (2) What qualifies a metaphysical relation as most apt for being called the '*F*-relation', where *F* names some concept? I claim that the reliability and success for causal learning of reliance on singular cues, and its independence of normative considerations, is good grounds for

assuming, for the sake of argument, that it is either inherently normative or not). For a good, opinionated discussion of the structure of our mature concept, see (Dinh & Danks, 2021, 2022).

[15] While in line with much of a similar proposal from White (2009a, 2009b, 2012a, 2012b, 2014), my argument does not require that all singular cues themselves derive from action originated by the agent.

[16] See also infant studies of causal perception. For example, (Leslie, 1982, 1984; Oakes, 1994).

positing the existence of a non-normative, singular causal relation. Further, I am happy to admit the existence of the kind of norm-loaded relation posited by a normative analysis. The locus of disagreement, then, is as to which earns the title, 'actual causation'. And here it is decidedly not the case that the nature of our proto-concept is beside the point. If my developmental proposal is correct, then at one stage in our development, we in fact do have a concept corresponding to the relation posited by a non-normative analysis of actual causation. Including this as part of the explanandum, a non-normative analysis arguably provides a more transparent and unified explanation than any the normative view could provide: the referent of our early concept is the same as that of our mature concept; just our means of identifying it evolve over time.

But such a tidy explanation is not available to the normative analysis. Early reliance on singular cues need not involve norms. It would be gratuitous to insist that nevertheless the concept has built-in normative content – even at this early stage.[17] But the proponent of a normative analysis must so insist. Alternatively, they could admit the early concept as non-normative. But doing so incurs a burden of explaining the relationship between the referent of our early concept (a non-normative singular relation) and our mature concept (an inherently normative relation that holds relative to a set of norms). One the non-normative analysis doesn't have. Once we adopt a holistic view of our concept of actual causation – one that includes its development – the normative analysis loses its advantage.

If we instead suppose that our mature concept is not inherently normative, then any advantage for the normative analysis is lost. If even our mature concept of actual causation is non-normative, then surely the metaphysical relation most apt to be called 'actual causation' is non-normative, as well.

§5 Conclusion

[17] There may be room to argue, in line with the suggestion of footnote 13, that norms do play a role at this early stage in their marking out certain singular cues as particularly salient. Studies would need to be run to check whether this is the case. The success of this argument would hinge on getting the right empirical results.

In sum, neither of the two principal reasons for inferring from the normativity of causal judgment to the normativity of causation itself pans out. A non-normative view can provide an equally adequate story for why, for whatever relation is proposed as an analysis of actual causation, that very particular relation is so qualified to go under the name 'actual causation'.

Statements and Declarations There are no competing interests to disclose.

References

Ahn, W., Kalish, C. W., Medin, D. L., & Gelman, S. A. (1995). The Role of Covariation Versus Mechanism Information in Causal Attribution. *Cognition*, *54*, 299–352.

Blanchard, T., & Schaffer, J. (2017). Cause Without Default. In H. Beebee, C. Hitchcock, & H. Price (Eds.), *Making a Difference: Essays on the Philosophy of Causation* (pp. 175–214). Oxford University Press. https://doi.org/10.1093/oso/9780198746911.003.0010

Bullock, M., Gelman, R., & Baillargeon, R. (1982). The Development of Causal Reasoning. In W. Friedman (Ed.), *The Developmental Psychology of Time* (pp. 209–254). Academic Press.

Burns, P., & McCormack, T. (2009). Temporal Information and Children's and Adults' Causal Inferences. *Thinking and Reasoning*, *15*, 167–196.

Campbell, J. (2020). *Causation in Psychology*. Harvard University Press.

Carey, S. (2009). *The Origin of Concepts*. Oxford University Press.

Danks, D. (2009). The Psychology of Causal Perception and Reasoning. In H. Beebee, C. Hitchcock, & P. Menzies (Eds.), *The Oxford Handbook of Causation*. Oxford University Press.

Dinh, P., & Danks, D. (2021). Causal Pluralism in Philosophy: Empirical Challenges and Alternative Proposals. *Philosophy of Science*, *88*(5), 761–772.

Dinh, P., & Danks, D. (2022). Causal Perception and Causal Inference: An Integrated Account. In P. Willemsen & A. Wiegmann (Eds.), *Advances in Experimental Philosophy of Causation* (pp. 81–100). Bloomsbury Academic.

Dretske, F. (1977). Referring to Events. *Midwest Studies in Philosophy*, *2*(1), 90–99.

Eells, E. (1988). Probabilistic Causal Interaction and Disjunctive Causal Factors. In J. H. Fetzer (Ed.), *Probability and Causality* (Vol. 192, pp. 189–209). Springer.

Eells, E. (1991). *Probabilistic Causality*. Cambridge University Press.

Gallow, J. D. (2021). A Model-Invariant Theory of Causation. *Philosophical Review*.

Hall, N. (2007). Structural equations and causation. *Philosophical Studies*, *132*(1), 109–136. https://doi.org/10.1007/s11098-006-9057-9

Halpern, J. (2008). Defaults and Normality in Causal Structures. *Proceedings of the Eleventh International Conference on Principles of Knowledge Representation and Reasoning*, 198–208.

Halpern, J., & Hitchcock, C. (2010). Actual Causation and the Art of Modeling. In *Causality, Probability, and Heuristics: A Tribute to Judea Pearl* (pp. 383–406). London: College Publications.

Halpern, J., & Hitchcock, C. (2015). Graded Causation and Defaults. *The British Journal for the Philosophy of Science*, *66*(2), 413–457.

Hart, H. L. A., & Honoré, T. (1985). *Causation in the Law*. Oxford University Press.

Hitchcock, C. (2001). A Tale of Two Effects. *Philosophical Review*, *110*, 361–396.

Hitchcock, C. (2017). Actual Causation: What's the Use? In *Making a Difference: Essays on the Philosophy of Causation* (pp. 116–131). Oxford University Press.

Hitchcock, C., & Knobe, J. (2009). Cause and Norm. *Journal of Philosophy*, *106*(11), 587–612.

Icard, T. F., Kominsky, J. F., & Knobe, J. (2017). Normality and Actual Causal Strength. *Cognition*, *161*, 80–93.

Kahneman, D., & Miller, D. (1986). Norm Theory: Comparing Reality to Its Alternatives. *Psychological Review*, *80*, 136–153.

Knobe, J. (2009). Folk Judgments of Causation. *Studies in History and Philosophy of Science Part A*, *40*(2), 238–242.

Knobe, J., & Fraser, B. (2008). Causal Judgment and Moral Judgment: Two Experiments. In W. Sinnott-Armstrong (Ed.), *Moral Psychology: The Cognitive Science of Morality* (Vol. 2, pp. 441–447). MIT Press.

Lagnado, D., & Sloman, S. (2004). The Advantage of Timely Intervention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *30*, 856–876.

Lagnado, D., & Sloman, S. (2006). Time as a Guide to Cause. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *32*(3), 451–460.

Leslie, A. (1982). The Perception of Causality in Infants. *Perception*, *11*, 173–186.

Leslie, A. (1984). Spatiotemporal Continuity and the Perception of Causality in Infants. *Perception*, *13*, 287–305.

Mascalzoni, E., Regolin, L., Vallortigara, G., & Simion, F. (2013). The Cradle of Causal Reasoning: Newborns' Preference for Physical Causality. *Developmental Science*, *16*(3), 327–335.

McCormack, T., Frosch, C., Patrick, F., & Lagnado, D. (2015). Temporal and Statistical Information in Causal Structure Learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *41*(2), 395–416.

McDonald, J. (2022). *Actual Causation: Apt Causal Models and Causal Relativism*. The Graduate Center, CUNY. http://academicworks.cuny.edu/gc_etds/4828

Mendelson, R., & Shulz, T. (1976). Covariation and Temporal Contiguity as Principles of Causal Inference in Young Children. *Journal of Experimental Child Psychology*, *22*, 408–412.

Menzies, P. (2004a). Causal Models, Token Causation, and Processes. *Philosophy of Science*, *71*(5), 820–832. https://doi.org/10.1086/425057

Menzies, P. (2004b). Difference Making in Context. In J. Collins, N. Hall, & L. A. Paul (Eds.), *Causation and Counterfactuals* (pp. 341–367). Oxford University Press.

Menzies, P. (2017). The Problem of Counterfactual Isomorphs. In H. Beebee, C. Hitchcock, & H. Price (Eds.), *Making a Difference: Essays on the Philosophy of Causation*. Oxford University Press.

Michotte, A. (1946). *The Perception of Causality*. Methuen & Co Ltd.

Oakes, L. (1994). Development of Infants' Use of Continuity Cues in Their Perception of Causality. *Developmental Psychology*, *30*(6), 869–879.

Paul, L. A., & Hall, N. (2013). *Causation: A User's Guide*. Oxford University Press.

Pearl, J. (2009). *Causality: Models, reasoning, and inference* (Second edition., 3rd printing..). Cambridge University Press. (Original work published 2000)

Pust, J. (2019). Intuition. *The Stanford Encyclopedia of Philosophy*, *Summer 2019*. <https://plato.stanford.edu/archives/sum2019/entries/intuition/>

Rips, L. (2011). Causation from Perception. *Perspectives on Psychological Science*, *6*(1), 77–97.

Rose, D., Sievers, E., & Nichols, S. (2021). Cause and Burn. *Cognition*, *207*.

Rottman, B. M., Kominsky, J. F., & Keil, F. C. (2014). Children use Temporal Cues to Learn Causal Directionality. *Cognitive Science*, *38*, 489–513.

Schaffer, J. (2012). Causal Contextualism. In M. Blaauw (Ed.), *Contrastivism in Philosophy: New Perspectives*. Routledge.

Shultz, T. R. (1982). Rules of Causal Attribution. *Monographs of the Society for Research in Child Development*, *47*, 1–51.

Shultz, T. R., Altmann, E., & Asselin, J. (1986). Judging Causal Priority. *British Journal of Developmental Psychology*, *4*, 67–74.

Shultz, T. R., Fisher, G. W., Pratt, C. C., & Rulf, S. (1986). Selection of Causal Rules. *Child Development*, *57*, 143–152.

Sytsma, J., & Livengood, J. (2021). Causal Attributions and the Trolley Problem. *Philosophical Psychology*, *34*(8), 1167–1191.

White, P. A. (1988). Causal Processing: Origins and Development. *Psychological Bulletin*, *104*, 36–52.

White, P. A. (1989). A Theory of Causal Processing. *British Journal of Psychology*, *80*, 431–454.

White, P. A. (1995). *The Understanding of Causation and the Production of Action: From Infancy to Adulthood*. Psychology Press: Taylor and Francis Group.

White, P. A. (2009a). Perception of Forces Exerted by Objects in Collision Events. *Psychological Review*, *116*, 580–601.

White, P. A. (2009b). Property Transmission: An Explanatory Account of the Role of Similarity Information in Causal Inference. *Psychological Bulletin*, *135*, 774–793.

White, P. A. (2012a). The Experience of Force: The Role of Haptic Experiences of Forces in Visual Perception of Object Motion and Interactions, Mental Simulation, and Motion-Related Judgments. *Psychological Bulletin*, *138*, 589–615.

White, P. A. (2012b). Visual Impressions of Pushing and Pulling: The Object Perceived as Causal is Not Always the One that Moves First. *Perception*, *41*, 1193–1217.

White, P. A. (2014). Singular Clues to Causality and Their Use in Human Causal Judgment. *Cognitive Science*, *38*, 38–75.

Willemsen, P., & Kirfel, L. (2019). Recent Empirical Work on the Relationship between Causal Judgments and Norms. *Philosophy Compass*, *14*.

Wolff, P. (2007). Representing Causation. *Journal of Experimental Psychology: General*, *136*, 82–111.

Wysocki, T. (forthcoming). Conjoined Cases. *Synthese*.

Yablo, S. (1992). Mental Causation. *The Philosophical Review*, *101*(2), 245–280. https://doi.org/10.2307/2185535

Young, M., & Sutherland, S. (2009). The Spatiotemporal Distinctiveness of Direct Causation. *Psychonomic Bulletin and Review*, *16*(4), 729–735.