Causal Perception and Causal Inference: An Integrated Account

David Danks[*] & Phuong (Phoebe) Ngoc Dinh

Carnegie Mellon University

5000 Forbes Ave

Pittsburgh, PA 15213

david@danks.org & pdinh@andrew.cmu.edu

Abstract

Causal learning is often thought to divide into two distinct types of cognitive processes and representations: causal perception and causal inference. In this chapter, we critically examine the evidence for two distinct kinds of cognitive processes, and show that extant experiments do not actually provide much evidence directly in favor of this pluralism. Research on causal learning has largely proceeded in two different paradigms, and so there are systematic methodological confounds that can explain the appearance of distinct processes. Moreover, the few experiments to investigate the relationship between causal perception and causal inference have provided suggestive evidence that they might be less distinct than commonly thought. We thus turn to the space of possible theories for (human) causal learning, and argue that there are natural theoretical options that have not yet been systematically explored. We describe one unexplored possibility in more detail—an integrated account based on inference to shared representations. In particular, this proposal holds that causal learners opportunistically use a wide range of features to infer the existence and strength of unobserved causal connections, and then explain, predict, and reason about the world around them on the basis of those inferred connections. We conclude by outlining key experiments to test the viability of this proposal.

[*] Corresponding author: david@danks.org

Keywords: Causal perception; Causal inference; Causal learning; Graphical models; Developmental trajectories

Index terms: Causal perception; Causal inference; Causal pluralism; Causal monism; Launching effect; Cognitive representations; Statistical cues

Much has been written about whether causation is a genuine feature of the world (for an overview, see Beebee, Hitchcock, & Menzies, 2009), but it seems undisputed that causation mediates much of human understanding and experience of the world. Causal knowledge enables us to predict future instances, explain past events, design interventions, categorize entities, reason about counterfactuals, and more (Sloman & Lagnado, 2015). As such, causation has been of longstanding interest to psychologists and philosophers alike. One prominent feature of causation is its variability: causal events exhibit diverse observable features, time courses, and necessary or sufficient conditions across domains. There is no single way that causation appears in our experience, no single phenomenological property that (seemingly) occurs for all causation. For example, imagine a rolling ball hitting a stationary ball, after which the latter immediately begins to move. Most adults would discern—easily, immediately, without prompting, and typically in one instance—that the former caused the latter to move (i.e., causal perception, Hubbard, 2013a, 2013b; Michotte, 1946/1963). In contrast, consider the case of determining whether smoking causes lung cancer. Here, learning appears effortful and explicit, requires multiple observations, and is amenable to top-down influences such as knowledge of other risk factors of lung cancer (i.e., causal inference, Cheng, 1997; Sloman & Lagnado, 2015).

In light of the seemingly distinct ways that humans learn and reason about causal events, as well as the diversity of events in the world that appear to involve causation, the idea of *causal pluralism*—the theory that there is a plurality of causal concepts and learning modes—has been repeatedly proposed in philosophy (Anscombe, 1971; Cartwright, 2004; Godfrey-Smith, 2010; Hall, 2004; Hitchcock, 2007, 2012; Psillos, 2010). Causal pluralism has also been suggested, though to a lesser extent, in psychology (Lombrozo, 2010; Waldmann & Mayrhofer, 2016). Even when psychologists do not explicitly endorse causal pluralism, they often tacitly assume it in their research programs: research on one kind of causal learning rarely incorporates insights from research on others, and efforts to investigate potential overlaps between (ostensibly different) causal concepts have largely been absent (with a few notable exceptions, such as Schlottmann & Shanks, 1992).

Despite the allure of causal pluralism, we contend that it stands on shaky empirical grounds. First, the apparent empirical distinction between different notions of causation discovered in psychological research is systematically confounded by methodological differences. Second, even if one concedes that there are truly different clusters of causal learning behaviors, there is research, albeit piecemeal and preliminary, to suggest that the corresponding conceptual boundary must be quite blurry. Even if causal pluralism provides a useful first approximation of human behavioral data, close inspection reveals the need for either substantial modifications or theoretical alternatives. Nonetheless, causal pluralism has remained a prominent view in both philosophy and psychology of causation, partly because extant monist theories all suffer from their own significant shortcomings.

In this chapter, we attempt to remedy this latter issue. We propose a new monist account of people's concept of causation, and provide a computational model of cognitive processes involving it. In particular, our monist account shows how a single kind of casual concept can nonetheless support multiple forms of causal learning and reasoning. This monist concept does not privilege any specific type of information a priori, but rather can be inferred from spatiotemporal, statistical, and mechanism information. We posit that causal learners opportunistically use any-and-all features to which they have epistemic access in order to infer causal connections, and then use those inferred connections to explain, predict, and reason about the world around them. We show that the new monist concept is capable of explaining existing empirical data on human causal learning, including data to suggest interactions between input of different modalities. We additionally aim to show that our proposal is not empirically vacuous, but rather makes novel predictions that have not previously been explored.

Before turning to our new account, though, we first survey the data that purportedly support causal pluralism. We then show that methodological confounds in experimental paradigms, measures, and explanatory foci undermine the conceptual boundaries proposed in causal pluralism. We also briefly discuss extant data that make a unitary causal concept seem plausible. We consider two existing monist proposals, each of which uses one of the proposed causal concepts to ground the other. We then introduce the basic tenets of our proposed monist causal theory, and provide a high-level explanation of how it can be computationally implemented. We show that this account makes testable predictions, and outline some preliminary investigations. We conclude with observations and lessons for both philosophy and psychology of causation as well as other domains.

A Brief Argument for Causal Pluralism

The most common form of causal pluralism in psychology posits two concepts and learning modes of causation: causal perception and causal inference. The first mode of learning—causal perception—is characteristically found in collisions or other direct physical causation. This mode hinges on signature perceptual features of dynamic events, such as the spatiotemporal contiguity between agents and recipients during launching (Michotte, 1946/1963; Yela, 1952), or the synchrony between the motion onsets of different objects in a chain of events (Hubbard & Ruppel, 2013; White & Milne, 1997). More recent research suggests that humans distinguish between some categories of causal interactions (e.g., launching vs. entraining) even in "low-level" vision (Kominsky & Scholl, 2020). Furthermore, causal perception appears impervious to top-down influences such as goals and prior knowledge, similar to some visual illusions (Firestone & Scholl, 2016). In most studies of causal perception, adults only need one exposure to determine the causal nature of the event. Notably, causal perception appears irresistible and phenomenologically salient (Michotte, 1946/1963) even if learners "know" otherwise given statistical dependency information (Schlottmann & Shanks, 1992). Infants develop the ability to perceive simple launches as causal between 6½ and 10 months of age (Leslie & Keeble, 1987; Oakes & Cohen, 1990), and as early as 4½ months of age with experience of self-generated action (Rakison & Krogh, 2012).

In contrast, the second mode of learning—causal inference—is characteristically found in learning from repeated experiences, as when one learns that aspirin relieves headaches (or red wine can produce them). Causal inference is sometimes subdivided into learning causal strength and learning causal structure, though these are not necessarily distinct cognitive processes (Griffiths & Tenenbaum, 2005). Spatiotemporal contiguity plays little-to-no role in causal inference. Instead, adults typically use contingency information between categorical variables (Rottman & Keil, 2012), covariation information between continuous factors (Marsh & Ahn, 2009; Soo & Rottman, 2018), deviations from base rates (Perales & Shanks, 2003), and other forms of statistical information. In experiments, participants extract statistical data from observation (Steyvers et al., 2003) or generate the data themselves (Hagmayer & Waldmann, 2007), and then draw conclusions about the existence, strength, and direction of causation. Causal inference is typically thought to be more effortful and explicit, and less phenomenologically salient, than causal perception. Top-down prior knowledge can readily guide causal inference (Hagmayer et al., 2011); for example, it can direct attention to causally relevant aspects of an event, or to potential interventions (Kushnir, Wellman, & Gelman, 2009). The earliest convincing evidence for children's causal inference was found in 19-month-old toddlers (Sobel & Kirkham, 2006). Causal inference becomes more sophisticated with development (McCormack et al., 2013; Waismeyer & Meltzoff, 2017).

Causal perception and causal inference present as strikingly different cognitive processes and behavioral patterns, and so some psychologists have argued that each learning mode requires a different kind of causal concept. For example, causal perception has been said to hinge on a perceptual concept in which causation is characterized by signature perceptual features that indicate a causal connection, such as spatiotemporal contiguity (White, 2014). Or causal perception may be grounded in a concept of causation as a mechanistic process that transfers power or a conserved quantity from one object to another in ways that yield perceivable signals (Wolff, 2014). These psychological proposals align nicely with process- or production-centric theories in philosophy, where causation is defined by either a conservation or invariance of some quantity through state changes (Dowe, 1992, 2000) or the propagation of causal influence through a chain of spatiotemporally contiguous events (Salmon, 1984, 1994).

In contrast, causal inference seems to be grounded in a concept of causation that emphasizes statistical information, interventions, and explicit prior knowledge. That is, causation for this type of learning is thought to be the statistical relations between causal variables (Cheng & Buehner, 2012; Tenenbaum et al., 2011) or post-intervention probabilities or counterfactuals (Sloman & Lagnado, 2005; Waldmann & Hagmayer, 2005). The statistical concept echoes several difference-making proposals in philosophy of causation, including those in which a cause is statistically correlated with its purported effect (Good, 1961a, 1961b), counterfactually related to the effect such that if it had not occurred then the effect would not have (counterfactual, Lewis, 1974), or manipulable to produce changes in the effect (interventionism, Menzies & Price, 1993; Woodward, 2005, 2011).

The ample (apparent) evidence for different types of causal learning, as well as different concepts and paradigmatic features, seems to support causal pluralism in psychology, one causation from causal perception and one from causal inference (for a more extensive synthesis, see Dinh, Danks, & Rakison, under review). This conclusion is reinforced by philosophical arguments for a similar position (Hall, 2004; Hitchcock, 2007). Causal pluralism has received explicit endorsement in the field (e.g., Lombrozo, 2010; or Waldmann & Mayrhofer, 2016 for a modified causal pluralism). Even when psychologists do not explicitly endorse causal pluralism, the field evolves as if causal perception and causal inference are indeed distinct clusters: many studies in causal inference exclude factors of causal inference that might be at play, and vice versa. Both in theory and practice, causal pluralism is arguably the default.

Methodological Challenges to Causal Pluralism

Although the main advance of this chapter is the proposal of a novel monist theory, we must first address the extensive body of work seemingly in support of causal pluralism. Our core response is that there is a natural alternative explanation (besides causal pluralism) for these behavioral and phenomenological data: namely, the systematic methodological differences between the two research areas. That is, we contend that the differences between causal perception and causal inference can potentially be explained by methodological confounds, rather than distinct concepts. Of course, this argument does not thereby establish causal monism, but by undermining the main argument in favor of causal pluralism, we open the door for consideration of novel theories. We focus here on three methodological confounds (but see Dinh & Danks, forthcoming for more systematic consideration of this challenge to causal pluralism).

One source of methodological divergence lies in the typical experimental paradigms for causal perception and causal inference. In causal perception research, participants usually judge individual events of bivariate causation (e.g., launching between two objects). The use of one-shot presentations in which the cause may or may not be efficacious means that events in causal perception paradigms often appear fully deterministic. In contrast, studies of causal inference usually encourage (if not require) that participants integrate data from multiple data points, whether presented as a table of summary statistics, a matrix of individual trials, or a sequential presentation of trials. Causal inference studies can involve both simple and complex causal relations, ranging from bivariate relations to causal webs with multiple causal mediators. Additionally, the causal relations studied in causal inference are often nondeterministic (with the probabilities provided through the statistical information conveyed to participants).

Another difference between causal perception and causal inference research is their measures. In causal perception research, adults typically answer questions about the power of the cause, whether through free-form responses, forced choices, or continuous rating scales. Alternately, implicit measures such as perception of overlap or expectations of the distance that the causal recipient should travel (i.e., representational momentum) are used in an effort to separate "low-level" causal percepts from the resulting "high-level," cognitively mediated inferences (Wagemans, van Lier, & Scholl, 2006). In contrast, measures in causal inference research span from predictions of future successes, to ratings of causal power on a continuous scale, to direct interventions on a causal system, and more. That is, measures of causal perception require that participants consider only the event they just watched (i.e., individual event, token), whereas those of causal inference often require that participants consider a group of trials or a causal type.

Yet another divergence between causal perception and causal inference research centers on the kinds of stimuli (and data) provided to participants. In causal perception, participants typically perceive dynamic events that unfold in space and time; temporal and spatial dimensions are explicitly presented in the stimuli. Other information about the kinematics of the event is also provided modally and directly (e.g., relative velocities, angle of approach). In contrast, the stimuli in causal inference research range widely in format, including dynamic events, diagrams of possible causal structures, static schematics of individual trials, or descriptions of events that accompany a summary table. In causal inference research, dynamic information such as space and time often need to be inferred, rather than being provided directly in the stimuli. It is even rarer for dynamic information to be presented modally in causal inference research.

The methods and measures of each research cluster are highly defensible if one starts with an appropriate understanding of the paradigmatic instances of causation for that research cluster. The methods and measures of causal perception research are tailored to phenomenological aspects of token events, while those of causal inference research are often cast at the level of causal types. As such, the theoretical perspectives that arise from one research cluster often struggle to account for information pertinent to the other research cluster. For example, many theories of causal perception have no formal account of how statistical information factors into the phenomenological salience of a launching event. Similarly, accounts of causal inference fumble at explaining or modeling how the experiential richness of causation arises from statistical data and top-down knowledge (cf. thick causation, Cartwright, 2004). To be clear, we are not suggesting that any causal perception or causal inference researchers (ourselves included!) have used incorrect or inappropriate methods. However, the significant differences in methods undermine our ability to draw strong inferences about differences in corresponding concepts, at least based on these empirical data. The argument for causal pluralism instead reduces to the intuitions with which we started this chapter. The methodologically splintered history of causal learning research makes it difficult to develop a principled investigation of the ways in which causal perception and causal inference might overlap.

The Empirical Case for a Unitary Account

Closer consideration of the empirical data actually provides some suggestive evidence that people might have a *unitary* concept of causation. In one direction, causal perception can exhibit traits typical of causal inference. For example, the perceptual triggers of experiences of causal perception, despite our previous characterizations, can be quite fuzzy and amenable to learning effects. With repeated exposure to delayed launches during experimental training, adults can perceive a delayed launch as causal, at least up to a point (Gruber, Fink, & Damm, 1957). Conversely, training with immediate launches narrows the temporal criterion for causal perception at test (Powesland, 1959). Adults misremember the order of event segments in a way that aligns with a causal interpretation if they perceived the entire event as causal (Bechlivanidis & Lagnado, 2016), and spatial criteria for causal perception are subject to similar effects: objects are judged as being spatially closer to each other when they are perceived as causally linked (Buehner & Humphreys, 2010). Relatedly, when a launch variant with spatial overlap between two objects is perceived as causal, adults underestimate the degree of overlap (Scholl & Nakayama, 2004). Interestingly, prior knowledge of typical features of causal actors (e.g., possession of dynamic parts, the ability to engage in self-propulsion) can constrain causal role assignment in a causal perception paradigm as early as 20 months of age (Rakison, 2006). These findings suggest that causal perception can be sensitive to top-down influences and training across trials, which are features more typically associated with causal inference.

In the other direction, causal inference is responsive to factors often associated with causal perception, particularly perceptual details about the dynamics of causal events. For example, young children struggle to discount misleading information about spatiotemporal contiguity when judging the outcome in a two-cause system. Compared to 9- and 10-year-olds, 5-year-olds were more likely to predict that an effect would occur immediately even when the mechanism was known to be slow (Schlottmann, 1999). Similarly, 3- to 3.5-year-olds were less likely to succeed at the blicket detector task when the objects hovered above the machine rather than placed on it (Kushnir & Gopnik, 2007). Even though adults can resist the allure of spatiotemporal contiguity in their causal judgments, they need clear reasons to do so, such as knowledge of a delayed mechanism (McGregor & Buehner, 2009). More recent research suggests that adults use small differences in time windows during causal inference to select between different potential causal structures (Bramley et al., 2018). Overall, perceptual information (e.g., space, time) seems to be integrated with other kinds of information during causal inference, and such details are assumed by learners even in the absence of explicit bottom-up information or direct instruction (Hagmayer & Waldmann, 2002).

Alternatives to Causal Pluralism: Two Grounding Accounts

Although causal pluralism has often functioned as the default position, alternative monist theories have started to emerge. Essentially all of these alternatives prioritize one of the two concepts (perceptual or statistical causality), and then explain the other concept in terms of the prioritized one. In one set of alternatives, the perceivable features of causation define the underlying concept and statistical features are based on that concept, so we refer to these monist accounts as *perceptual grounding theories*. When we look at causal perception, we find that cues such as spatiotemporal contiguity and self-propulsion signal agent and recipient roles in causal interactions, and thereby license the inference to a causal relation (Rips, 2011; White, 2014). Perceptual grounding theories often propose that learners become sensitive to these cues through lifelong experience with causal events that begin in their own experience of exerting change on their surroundings. Alternatively, causation might correspond to a continuous chain of events of the right, force-transmitting kind or mechanism (Ahn & Kalish, 2000; Wolff, 2014). Or learners may hold intuitive theories about momentum and physical causality that are made imprecise by perceptual noise, other prior beliefs, or even assumed uncertainty (Gerstenberg & Tenenbaum, 2017). Regardless of the exact story, all perceptual grounding accounts give perceptual features content beyond their perceivability, and thereby attempt to explain learning in contexts previously thought to be outside of the scope of causal perception. For example, perceptual grounding accounts have been offered to explain counterfactual simulations of launching (Gerstenberg, Halpern, & Tenenbaum, 2015) or the (context-appropriate) downweighting of spatiotemporal contiguity during causal judgments in cognitive development (Schlottmann, 1999).

Perceptual grounding accounts have some intuitive appeal and potential explanatory power, but also have a number of open questions. For example, even an extended or generalized concept of perceptual causation does not seem capable of representing events with few or no immediately perceivable features, such as the causal connection between antidepressants and depressive symptoms. This causal relation has a noisy time course, many mediators or defeaters, and no clear perceptual signatures for the learner to use. Many other phenomena also escape ordinary human perception, and yet we clearly learn causal relations, such as the discovery of the general shape of planetary orbits and what forces govern such a shape. Perhaps most importantly, most accounts of perceptual grounding have no straightforward way of differentiating observation from intervention even though human learners do (Waldmann & Hagmayer, 2005).

In a different set of monist alternatives, researchers aim to prioritize a concept of causation based on statistical information, and then build other causal concepts (e.g., perceptual causality) on top of that concept (Kemp, Goodman, & Tenenbaum, 2010), and so we refer to these as *statistical grounding theories*. These statistical features can manifest as correlations (Good, 1961a, 1961b), contingencies (Perales & Shanks, 2003), interventions (Gebharter, 2017; Woodward, 2005), or counterfactuals (Lewis, 1974). Regardless of the types of statistical information represented in a causal concept, learners can use them to learn and reason about both token- and type-level claims: Smoking increases the chances of developing lung cancer at both the population and individual levels. On these accounts, the perceptual features of a causal event are only salient (if at all) in virtue of their statistics (Woodward, 2011). For example, spatiotemporal contiguity almost always predicts successful launching (given certain conditions e.g., the recipient is not too heavy for the agent), and so the reliance on perceptual cues is entirely reducible to the use of highly statistically significant cues. Notably, several statistical grounding accounts can differentiate between observation and intervention (e.g., the *do* operator in causal graphical models).

Despite their explanatory power, statistical grounding accounts also face open questions and challenges. First, these theories struggle with the phenomenological salience and richness of many daily causal experiences. There is a clear phenomenological difference between the causal perception of our dog chasing after the neighbor's cat and calculations (even if implicit) of the probability of the cat getting hurt. The perception and corresponding beliefs are richer than the statistical properties that underlie them, and so suggest that statistical grounding theories must provide additional explanations to account for the richness of causal perception. Causal perception can provide compelling impressions that diverge from (and cannot be extinguished by) statistically driven conclusions, even if those impressions do not dominate learning outcomes in the end (Schlottmann & Shanks, 1992). Statistical grounding theories also need a way to represent mechanistic information, as there is empirical evidence that people do not conceive of causal mechanisms in purely statistical terms (Ahn & Kalish, 2000). And although causal perception is not completely immune to top-down influences, it does seem to be significantly more resistant to such effects than explicitly statistical causal beliefs. As with perceptual grounding theories, these open questions do not thereby show that these theories are false, but they should temper our potential enthusiasm for these avenues towards a monist theory.

Causal Monism: A New Alternative to Causal Pluralism

Causal pluralism faces significant empirical challenges, and there are legitimate theoretical worries about the monist causal learning theories that privilege either causal perception or causal inference. Given these concerns, we develop a different type of monist theory of causal learning and reasoning in this section. We present this theory below, but we emphasize that our primary goal here is to broaden the theoretical space to include a novel, empirically testable possibility. Systematic experiments remain a subject for future work. While this theory may ultimately be empirically falsified, our understanding of the potential relationships between causal perception and causal inference is, we suggest, significantly advanced by our proposed theory.

At a high level, this monist theory posits that people have a single, relatively amodal representation of "unobservable causal connection." This theory further posits that people are opportunists: they use any-and-all clues accessible to them in their efforts to infer these unobservable causal connections, including perceptual cues, statistical information, and verbal instructions. Given an inferred causal connection, one can then reverse the information flow to predict or infer other, not-yet-observed clues. For example, if I infer a causal connection on the basis of statistical information, then I can thereby reason that there is probably some measure of spatial and/or temporal contiguity (which might not be observable) mediating that connection. Importantly, these different sources of evidence (and targets for prediction and reasoning) do not correspond to different underlying concepts of causation, but are simply different pathways towards a single concept. This is analogous to the way that my concept of DOG can be activated by an image of a dog, the sound of a bark, or someone telling me about their dog (though see, e.g., Barsalou, 1999 or Machery, 2009 for arguments that no single, shared concept is activated in these different cases). That is, we are not proposing a monist "cluster concept" of causation where "C causes E" sometimes means *X* and sometimes means *Y*. Rather, this theory posits that there is a single coherent (overarching) concept of causal connection that underlies our representations of causal structure in the world, but we use many different types of information to infer its existence in particular token cases or types of events.

This high-level proposal can be made precise by formalizing the theory using graphical models. Importantly, we are here using graphical models to represent informational relations, not necessarily causal ones; there is no problematic circularity of people (implicitly) assuming that causal connections *cause* statistical patterns or spatiotemporally contiguous change. The graphical model instead represents the relevant evidential and informational relations. (Of course, as with any graphical model, inferences can occur in both directions along an edge; the directions provide the pattern of informational dependencies, not restrictions on permissible inferences.) We start by considering a single case or event, where that case might include any of the following information about two factors C and E:[1]

1. Spatiotemporally contiguous change in C and E (e.g., a launching event or linked changes in magnitude)
2. Co-occurrence of C and E (e.g., both are present or both are absent)
3. Control (e.g., via intervention) of one or both of C and E
4. Verbal content directly about C and/or E
5. Background experiences that determine prior expectations (discussed more below)

We can graphically represent the inferential relations as in Figure 1. Specifically, there is an unobserved variable "C is causally connected to E" (henceforth, *CC*) whose possible multidimensional values encode presence/absence, causal strength, and any mechanistic or spatiotemporal details. *CC* is an informational driver of the possibly-observed factors corresponding to the elements in this list.

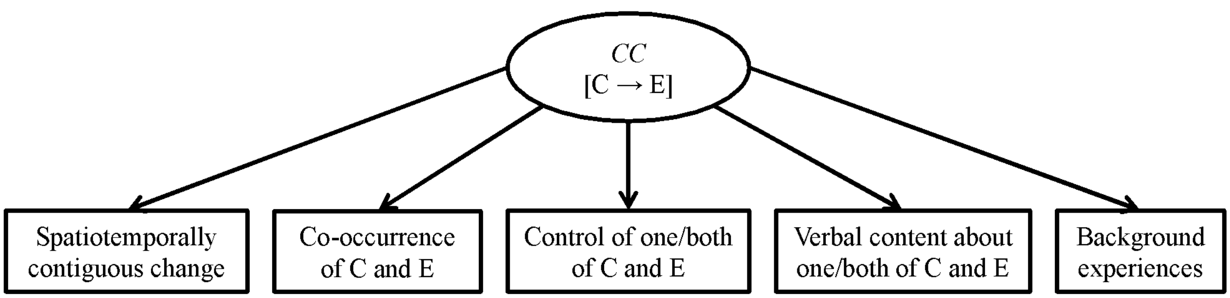

Figure 1: Graphical model for monist inference

Full specification of the (informational) graphical model requires the relevant likelihood functions of each observable factor given the presence/absence and strength of an unobserved causal connection. Importantly, the graphical model structure implies that we do not need to consider direct factor-factor informational relations; everything is mediated through the unobserved causal connection. For example, the full graphical model includes *P*(co-occurrence | *CC*), but not *P*(co-occurrence | verbal content). Put in different language, the state of *CC*

[1] This list is not necessarily complete and exhaustive. One advantage of this monist account is that it is readily extensible if we discover that people employ some other type of cue or feature in causal learning: we simply add that type of feature as another "leaf node" that can be used opportunistically to infer an unobservable causal connection.

"screens off" the different potential sources of causal information. As such, inferences to *CC* given some input enable subsequent learning and prediction of new input, whether of the same type or of a different type. The inference to *CC* need not be empirically accurate by some scientific standard; learners may infer *CC* mistakenly or hypothetically, given the data they observe. Most critically, inferences to the unobserved (or unobservable) variable *CC* are not an add-on feature, but rather are the unifying feature that binds together various instances of causal learning (similarly to the suggestion that abstraction to unobserved relations is key to human cognition, as argued by Penn, Holyoak, & Povinelli, 2008; Penn & Povinelli, 2007).

Given a single case or event, one can straightforwardly infer the probability of a causal connection of various types and strengths. Mathematically, this inference is simply standard Bayesian updating (though without requiring any of the usual metaphysical or normative baggage of Bayesianism). Some of the observable factors might not actually be observed, in which case they play no role in the inference. Everything that is observed, however, factors into the update about the probabilities of different values for *CC*. No type of observable evidence is privileged on theoretical or conceptual grounds for this inference, though they might have different informational values. If the probability distribution over *CC* changes due to the inferential update, then the informational impact of that change can "flow" back to the other factors (i.e., the ones that are observable, but not actually observed) in the usual way. That is, if I infer that there is probably a causal connection in this case because of factor *A*, then that change in my belief can also produce a revised expectation about factor *B* in this case (even if I have not yet observed *B*).

If we observe multiple cases of a type, then knowledge about the probabilities for *CC* can be transferred forward in standard ways.[2] Essentially, we are learning the "base rates" of causal connections with particular properties in instances of this type. Alternately, we might be explicitly told about type-level causal connections, which thereby set the relevant prior probabilities. Type-level information about different types could also exhibit structure; for example, if I know that dogs typically have a particular causal connection, then I might reasonably update my prior probabilities about this causal connection in wolves. These

[2] Multiple cases of a type can be represented mathematically using plate notation to capture the idea that these are all instances of the same class.

inferences can readily be captured through additional structure linking various *CC* variables (similarly to graphical model representations of between-concept relevance relations, as in Danks, 2014). The long-term result is that we can learn both whether a causal connection exists in a particular case (e.g., does this aspirin relieve this headache?) and also whether there are type-level causal connections (e.g., does aspirin generally relieve headaches?).

This sketch of a computational model can easily be made fully precise; none of the components or pieces are particularly esoteric or unusual (at least, within computational modeling). We have omitted the mathematics simply to help maintain focus on the conceptual and theoretical aspects of the proposal. That being said, we do need to provide here an account of the source of the likelihood functions, even if we skip the precise mathematical formulation. Adult humans clearly have various domain-dependent expectations about how causal connections manifest in the world, whether in terms of spatiotemporal, statistical, or manipulation relationships. For example, adults expect physical causation to be largely deterministic; psychological causation, not so much (Yeung & Griffiths, 2015; Strickland, Silver, & Keil, 2017). As another example, people plausibly expect that a causal connection in the realm of physical objects should (if it exists) exhibit spatiotemporally contiguous changes, while a causal connection in the realm of macroscale biological mechanisms should (if it exists) exhibit significant temporal gaps between the cause and effect. And these expectations express exactly the information required for the likelihood functions. That is, people's intuitive (default) expectations about how causal connections might manifest in observable properties are precisely the relevant likelihoods (see Danks, 2018 for a complementary account of default expectations for statistical information).

This monist theory can straightforwardly explain the empirical findings in causal perception and causal inference. Given spatiotemporal information, the learner infers the possible existence of an unobserved causal connection in this token, and then uses that updated belief to provide ratings or answer other questions about the causal connection. Given covariation information summed over multiple instances, the learner infers the possible existence of a pattern of causal connections for the type, and then provides strength ratings, determines post-intervention probabilities, and so forth. In each case, the relevant cognitive processes are simply part of the inferences posited by the monist theory. Of course, the monist theory also predicts that people will draw a number of conclusions from these inferences that are not implied by either causal perception or causal inference in isolation; we return to those predictions shortly. The key observation here is that this monist theory subsumes the existing theories in the special case where our experimental methodologies consider only perceptual or only statistical inputs. The monist theory thus fits cleanly with our earlier critique of the evidence for pluralism, as it contends that the lack of (apparent) interaction between causal perception and causal inference is principally because we have not systematically looked for it (though recall the studies outlined earlier in “The Empirical Case for a Unitary Account”).

This last observation about the monist theory suggests a potential fatal flaw: the theory might appear to be simply the concatenation of (the mathematics of) causal perception plus (the mathematics of) causal inference. Almost any set of non-contradictory theories can be “unified” by concatenating representations of them in a common mathematical language, but such a “unification” would tell us essentially nothing about the actual nature of the mind. There must be some content to the monist theory that goes beyond the union of the subsumed theories, else we have not provided anything more than a mathematical parlor trick. We thus conclude this section by outlining three key empirical predictions that are, to the best of our knowledge, distinctive to this monist account (and not implied by the concatenation of causal perception and causal inference).

First, people should be able to make predictions symmetrically from one type of information to another. For example, when shown a canonically causal perception-evoking event, people should be able to make predictions about post-intervention probabilities for this type of event. Or when provided with covariational data, people should make some inferences about the likelihood of spatiotemporally realized mechanisms in various tokens (perhaps different mechanisms in different tokens). The standard pluralist picture implies that this type of cross-task information transfer should be difficult or noisy. Monist theories that ground perception in inference (or vice versa) predict that cross-task transfer should be asymmetric: easy when going from perceptual information to statistical, and hard in the other direction (or the opposite prediction if inference is grounded in perception). In contrast, our monist account proposes that this cross-task transfer should be straightforward and symmetric.

Second, information about one type of feature should inform future learning using other features. Probabilistic updates of *CC* not only inform reasoning and prediction, but also influence future learning. For example, recall that people judge objects to be closer together when they are believed to be causally related (Buehner & Humphreys, 2010). This effect was demonstrated entirely within the domain of causal perception (i.e., perceptions following actual launching or delayed launching events). This monist theory implies that a similar perceptual effect should occur if people are instead provided with covariational information indicating that a causal connection is almost certainly present in this token, even if people never see the actual collision event. There are clearly methodological challenges in testing this prediction; for example, one would want to minimize, or at least measure for later statistical control, any potential (perceptual) mental simulation by participants in response to the statistical information. Nonetheless, this prediction is distinctive to this monist theory.

Third, both of these effects—cross-task/evidence transfer in learning and in reasoning—are predicted to be entirely (or mostly) mediated by *CC*. In our proposed account, the information transfer happens because of updates to aspects of *CC*, rather than direct inference from one kind of observable content to another. This latter possibility is what one would expect if there were learned associations between different types of observable information. There are clearly correlations between the different observable signals of causation, so one might attempt to explain cross-task transfer in terms of direct learned associations between those signals. For example, spatiotemporally contiguous change might be directly associated with intervention counterfactuals since those two go together frequently in a learner's experience. In contrast, our proposed theory implies that this information flow is via *CC*, and so fixing the values of *CC* (either statistically or causally) should block the cross-task transfer. We are currently developing an experimental design to test this more subtle prediction.

We conclude our proposal by describing some initial data that suggest that perceptual and statistical cues to causation can interact flexibly to influence the outcome of causal learning. We describe only the qualitative phenomena in this chapter for space reasons.[3] This experiment was designed to challenge a version of causal pluralism in which the multiple concepts are relatively independent. As a result, the experiment cannot distinguish between monism and an interactive

[3] The raw data for the experiments described here can be found at https://osf.io/263x5/.

version of causal pluralism whereby different causal concepts can strongly influence and constrain one another during learning (Waldmann & Mayrhofer, 2016).[4] (For attempts at developing differential predictions for these latter two theoretical possibilities, see Dinh, Danks, & Rakison, under review.) Nonetheless, the results point towards exactly the type of symmetry and flexibility between cue-types that is predicted primarily (though perhaps not exclusively) by our monist account. These experiments were all conducted with adults on Amazon Mechanical Turk, and all pitched perceptual and statistical cues against one another in a series of 12 animated, dynamic events. All events involved the motion onset of a stationary object, in response to either contact with a moving object (Launch event) or a rapid series of color changes between pink and purple in the stationary object before settling on purple (Blink event). These two types of events were chosen because we found a sharp divergence between expectations after a single successful instance: Launches were expected to be strong and reliable causes of motion, whereas Blinks were expected to be weak and unreliable.

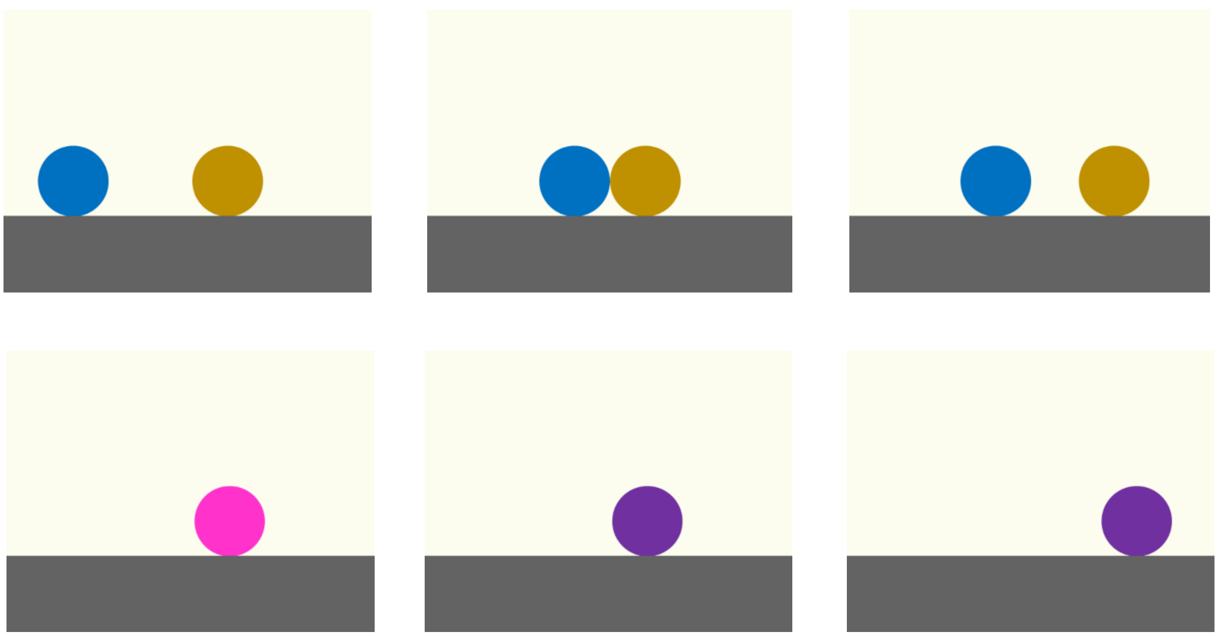

Figure 2: Screenshots of successful Launch (top row) and successful Blink (bottom row) events. In unsuccessful events, the recipient remained stationary.

The key experimental conditions contrasted these strong "perceptual" expectations against strong statistical evidence. For example, a participant might see a sequence of Launch

[4] We thank Nick Ichien for this important observation.

events that were only 75% successful, or she might see a sequence of Blink events that were 100% successful. According to standard causal pluralism and perceptual grounding theories, participants should give relatively constant cumulative judgments across a sequence of Launch events, regardless of statistical information (since the perceptual cues should be overwhelming and automatic). In contrast, Blink events do not trigger causal perception, and so learning should proceed according to causal inference (and thereby produce a standard learning curve). Statistical grounding theories arguably predict that the statistical information should play a major role, though perhaps significantly attenuated in Launch sequences due to the strong prior expectation of a deterministic relationship. And our monist theory predicts a complex pattern of learning (details omitted) as the perceptual and statistical cues are all used opportunistically to learn about the type-level causal connections.

Results bore out the predictions of our monist theory with high congruence between our two different measures. The first question asked participants to rate the extent to which an apparent cause (Launch or Blink) made the stationary ball move (-100 to +100). The second question asked participants to estimate the number of cases in which they would expect the stationary ball to move, given 100 cases in which the factor of that series (Launch or Blink) was present. For both questions, participants were asked to consider all events that they had seen in that series. When every event was successful (i.e., if determinism held), Launch sequences led to causal perception-like behavior but Blink sequences led to causal inference-like behavior. (And cumulative causal judgments for Blink sequences never reached the levels of Launch sequences.) When the sequence was nondeterministic, participant behavior was more complex than predicted by either pluralism or the grounding theories. If the first failure event happened on the very first trial, then participant judgments were low from the outset, regardless of whether it was a Launch or Blink event. A single failure at the very start was sufficient to largely eliminate the exclusive use of spatiotemporal cues. And if the first failure event occurred on a later trial, then participant judgments for Launch sequences dropped significantly more than for Blink sequences after that first failure. That is, the statistical information of a single failure (regardless of location in the sequence) had a significantly larger impact on the Launch sequence judgments, but those were supposed to be the judgments based on causal perception and so more resistant to statistical information! Similar findings occurred across a range of variations in timing of the first failure event, as well as pattern of failure events across the sequence. Our tentative conclusion is that people do not a priori privilege one type of information in causal learning. Rather, people opportunistically use and integrate information from diverse cues to infer aspects of the unobserved causal structure underlying their observations.

Conclusion

In this chapter, we briefly presented the central tenets and empirical support for causal pluralism, which is the proposal that human causal learning relies on two distinct kinds of causal concepts and modes of learning: causal perception and causal inference. We discussed methodological differences between the two research clusters and showed that they confound with claims of conceptual distinction between causal perception and causal inference. One alternative if we reject causal pluralism is to reduce one form of causal learning to the other, and so we reviewed two grounding accounts—perceptual grounding and statistical grounding. In contrast, we have proposed a third alternative, namely a monist account in which learners may use any and all cues accessible to them to infer the existence of an unobserved causal relation represented with a single, amodal causal concept. And once learners infer such causal connections, they can then make predictions about future occurrences or inferences about other, not-yet-observed, or unobservable features of the inferred relation. Finally, we provided three testable predictions of the monist account to distinguish it from a mere (mathematical) concatenation of existing theories. It remains to be seen whether our predictions will bear out with future empirical tests. Most importantly, we wish to call attention to a lack of systematic investigation into the ways in which causal perception and causal inference interact, despite extant results pointing to that possibility.

This chapter exemplifies an additional exercise relevant to philosophers and psychologists of causation and beyond. Any theory of empirical import must rely on findings that are discovered through particular research paradigms and methods. At the same time, theories provide the assumptions and constructs that guide and constrain their own experimental paradigms and methods. The interplay between theory and method can be a virtue: for example, theory-mediated measurements enable the precise quantification of key parameters relevant to a theory when done intentionally (Harper, 2007). Yet without that intentional design, methodological specifics can confound with and undermine inferences from data. We suspect this has been the case with research in causal learning. We recommend the incorporation of more diverse methods and measures in future research of causal concepts, as well as an investigation of the assumptions underlying those methods and how they might bear out in the data.

Acknowledgments

Thanks to the editors, Nick Ichien, and David Rakison for valuable comments and feedback on an earlier draft of this chapter.

References

Ahn, W. K., & Kalish, C. W. (2000). The role of mechanism beliefs in causal reasoning. In R. Wilson & F. Keil (Eds.), *Explanation and Cognition* (pp. 199-225). Cambridge, MA: MIT Press.

Anscombe, G. E. M. (1971). Causality and Determination. Reprinted in E. Sosa and M. Tooley (Eds.), *Causation*, 1993 (pp. 88-104). Oxford: Oxford University Press.

Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences, 22*, 577-660.

Bechlivanidis, C., & Lagnado, D. A. (2016). Time reordered: Causal perception guides the interpretation of temporal order. *Cognition*, *146*, 58-66. https://doi.org/10.1016/j.cognition.2015.09.001

Beebee, H., Hitchcock, C., & Menzies, P. (Eds). (2009). *The Oxford Handbook of Causation*. Oxford University Press.

Bramley, N. R., Gerstenberg, T., Mayrhofer, R., & Lagnado, D. A. (2018). Time in causal structure learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 44*(12), 1880-1910. https://doi.org/10.1037/xlm0000548

Buehner, M. J., & Humphreys, G. R. (2010). Causal contraction: Spatial binding in the perception of collision events. *Psychological Science*, *21*(1), 44-48. https://doi.org/10.1177/0956797609354735

Buehner, M. J., & McGregor, S. J. (2009). Contingency and contiguity trade-offs in causal induction. *International Journal of Comparative Psychology*, *22*(1), 19-42. Retrieved from https://escholarship.org/uc/item/8tb8w6f1.

Cartwright, N. (2004). Causation: One word, many things. *Philosophy of Science*, *71*(5), 805-819. https://doi.org/10.1086/426771

Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review, 104*(2), 367-405. https://doi.org/10.1037/0033-295X.104.2.367

Cheng, P. W. & Buehner, M. J. (2012). Causal Learning. In K. J. Holyoak & R. G. Morrison (Eds.),The Oxford Handbook of Thinking and Reasoning (pp. 210-233). Oxford, UK: Oxford University Press.

Danks, D. (2014). *Unifying the mind: Cognitive representations as graphical models*. Cambridge, MA: The MIT Press.

Danks, D. (2018). Privileged (default) causal cognition: A mathematical analysis. *Frontiers in Psychology, 9*: 498. https://doi.org/10.3389/fpsyg.2018.00498

Dinh, P., & Danks, D. (forthcoming, 2021). Causal pluralism in philosophy: Empirical challenges and alternative proposals. *Philosophy of Science, 88*(5).

Dinh, P. N., Danks, D., & Rakison, D. H. (under review). Causal Perception and Causal Inference: A Methodological and Theoretical Synthesis. Department of Psychology, Carnegie Mellon University.

Dowe, P. (1992). An empiricist defence of the causal account of explanation. *International Studies in the Philosophy of Science*, *6*(2), 123-128. https://doi.org/10.1080/02698599208573420

Dowe, P. (2000). *Physical Causation*. Cambridge, UK: Cambridge University Press.

Firestone, C., & Scholl, B. J. (2016). Cognition does not affect perception: Evaluating the evidence for “top-down” effects. *Behavioral and Brain Sciences*, *39*, e229. https://doi.org/10.1017/S0140525X15000965

Gebharter, A. (2017). Causal exclusion and causal Bayes nets. *Philosophy and Phenomenological Research*, *95*(2), 353-375. https://doi.org/10.1111/phpr.12247

Gerstenberg, T., Halpern, J. Y., & Tenenbaum, J. B. (2015). Responsibility judgments in voting scenarios. In D. C. Noelle et al. (Eds.), *Proceedings of the 37th Annual Conference of the Cognitive Science Society* (pp. 788-793). Austin, TX: Cognitive Science Society.

Gerstenberg, T., & Tenenbaum, J. B. (2017). Intuitive theories. In M. R. Waldmann (Ed.), *Oxford Handbook of Causal Reasoning*, 515-548. Oxford, UK: Oxford University Press.

Godfrey-Smith, P. (2010). Causal pluralism. In H. Beebee, C. Hitchcock, and P. Menzies (Eds.), *Oxford Handbook of Causation* (pp. 326–337). Oxford: Oxford University Press. https://doi.org/10.1093/oxfordhb/9780199279739.003.0017

Good, I. J. (1961a). A causal calculus (I). *The British Journal for the Philosophy of Science*, *11*(44), 305-318.

Good, I. J. (1961b). A causal calculus (II). *The British Journal for the Philosophy of Science*, *12*(45), 43-51.

Griffiths, T. L., & Tenenbaum, J. B. (2005). Structure and strength in causal induction. *Cognitive Psychology*, *51*(4), 334-384. https://doi.org/10.1016/j.cogpsych.2005.05.004

Gruber, H. E., Fink, C. D., & Damm, V. (1957). Effects of experience on perception of causality. *Journal of Experimental Psychology, 53*(2), 89-93. https://doi.org/10.1037/h0048506

Hagmayer, Y., Meder, B., von Sydow, M., & Waldmann, M. R. (2011). Category transfer in sequential causal learning: The unbroken mechanism hypothesis. *Cognitive Science*, *35*(5), 842-873. https://doi.org/10.1111/j.1551-6709.2011.01179.x

Hagmayer, Y., & Waldmann, M. R. (2002). How temporal assumptions influence causal judgments. *Memory & Cognition*, *30*(7), 1128-1137. https://doi.org/10.3758/BF03194330

Hagmayer, Y., & Waldmann, M. R. (2007). Inferences about unobserved causes in human contingency learning. *Quarterly Journal of Experimental Psychology*, *60*(3), 330-355. https://doi.org/10.1080/17470210601002470

Hall, N. (2004). Two concepts of causation. In J. D. Collins, E. J. Hall, & L. A. Paul (Eds.), *Causation and Counterfactuals* (pp. 225–276). Cambridge, Massachusetts: MIT Press.

Harper, W. (2007). Newton's methodology and Mercury's perihelion before and after Einstein. *Philosophy of Science*, *74*(5), 932-942. https://doi.org/10.1086/525634

Hitchcock, C. (2007). How to be a causal pluralist. In P. Machamer & G. Wolters (Eds.), *Thinking about causes: From Greek philosophy to modern physics* (pp. 200-221). Pittsburgh, PA: University of Pittsburgh Press.

Hitchcock, C. (2012). Portable causal dependence: A tale of consilience. *Philosophy of Science*, *79*(5), 942-951. https://doi.org/10.1086/667899

Hubbard, T. L. (2013a). Phenomenal causality I: Varieties and variables. *Axiomathes*, *23*(1), 1-42. https://doi.org/10.1007/s10516-012-9198-8

Hubbard, T. L. (2013b). Phenomenal causality II: Integration and implication. *Axiomathes*, *23*(3), 485-524. https://doi.org/10.1007/s10516-012-9200-5

Hubbard, T. L., & Ruppel, S. E. (2013). Ratings of causality and force in launching and shattering. *Visual Cognition*, *21*(8), 987-1009. https://doi.org/10.1080/13506285.2013.847883

Kemp, C., Goodman, N. D., & Tenenbaum, J. B. (2010). Learning to learn causal models. *Cognitive Science*, *34*(7), 1185-1243. https://doi.org/10.1111/j.1551-6709.2010.01128.x

Kominsky, J. F., & Scholl, B. J. (2020). Retinotopic adaptation reveals distinct categories of causal perception. *Cognition*, *203*, 104339. https://doi.org/10.1016/j.cognition.2020.104339

Kushnir, T., & Gopnik, A. (2007). Conditional probability versus spatial contiguity in causal learning: Preschoolers use new contingency evidence to overcome prior spatial assumptions. *Developmental Psychology, 43*(1), 186-196. https://doi.org/10.1037/0012-1649.43.1.186

Kushnir, T., Wellman, H. M., & Gelman, S. A. (2009). A self-agency bias in preschoolers' causal inferences. *Developmental Psychology, 45*(2), 597–603. https://doi.org/10.1037/a0014727

Legare, C. H., Gelman, S. A., & Wellman, H. M. (2010). Inconsistency with prior knowledge triggers children's causal explanatory reasoning. *Child development*, *81*(3), 929-944. https://doi.org/10.1111/j.1467-8624.2010.01443.x

Leslie, A. M., & Keeble, S. (1987). Do six-month-old infants perceive causality?. *Cognition*, *25*(3), 265-288. https://doi.org/10.1016/S0010-0277(87)80006-9

Lewis, D. (1974). Causation. *The Journal of Philosophy*, *70*(17), 556-567. https://doi.org/10.2307/2025310

Lombrozo, T. (2010). Causal-explanatory pluralism: How intentions, functions, and mechanisms influence causal ascriptions. *Cognitive Psychology*, *61*(4), 303-332. https://doi.org/10.1016/j.cogpsych.2010.05.002

Machery, E. (2009). *Doing without concepts*. New York: Oxford University Press.

Marsh, J. K., & Ahn, W. K. (2009). Spontaneous assimilation of continuous values and temporal information in causal induction. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *35*(2), 334–352. https://doi.org/10.1037/a0014929

McCormack, T., Simms, V., McGourty, J., & Beckers, T. (2013). Blocking in children's causal learning depends on working memory and reasoning abilities. *Journal of Experimental Child Psychology*, *115*(3), 562-569. https://doi.org/10.1016/j.jecp.2012.11.016

Menzies, P., & Price, H. (1993). Causation as a secondary quality. *The British Journal for the Philosophy of Science*, *44*(2), 187-203. https://doi.org/10.1093/bjps/44.2.187

Michotte, A. (1946). La perception de la causalité. Louvain: Institut Superior de Philosophie, 1946. English translation of updated edition by T. Miles & E. Miles, *The Perception of Causality*, Basic Books, 1963.

Oakes, L. M., & Cohen, L. B. (1990). Infant perception of a causal event. *Cognitive Development*, *5*(2), 193-207. https://doi.org/10.1016/0885-2014(90)90026-P

Penn, D. C., Holyoak, K. J., & Povinelli, D. J. (2008). Darwin's mistake: explaining the discontinuity between human and nonhuman minds. *Behavioral and Brain Sciences*, *31*(2), 109-178. https://doi.org/10.1017/S0140525X08003543

Penn, D. C., & Povinelli, D. J. (2007). Causal cognition in human and nonhuman animals: A comparative, critical review. *Annual Review of Psychology*, *58*, 97-118. https://doi.org/10.1146/annurev.psych.58.110405.085555

Perales, J. C., & Shanks, D. R. (2003). Normative and Descriptive Accounts of the Influence of Power and Contingency on Causal Judgement. *The Quarterly Journal of Experimental Psychology Section A*, *56*(6), 977–1007. https://doi.org/10.1080/02724980244000738

Powesland, P. F. (1959). The effect of practice upon the perception of causality. *Canadian Journal of Psychology/Revue Canadienne de Psychologie, 13*(3), 155-168. https://doi.org/10.1037/h0083773

Psillos, S. (2010). Causal pluralism. In R. Vanderbeeken & B. D'Hooghe (Eds.), *Worldviews, science, and us: Studies of analytic metaphysics: A selection of topics from a methodological perspective*. World Science Publishers. https://doi.org/10.1142/9789814299053_0009

Rakison, D. H. (2006). Make the first move: How infants learn about self-propelled objects. *Developmental Psychology, 42*(5), 900–912. https://doi.org/10.1037/0012-1649.42.5.900

Rakison, D. H., & Krogh, L. (2012). Does causal action facilitate causal perception in infants younger than 6 months of age?. *Developmental Science, 15*(1), 43-53. https://doi.org/10.1111/j.1467-7687.2011.01096.x

Rips, L. J. (2011). Causation from perception. *Perspectives on Psychological Science*, *6*(1), 77-97. https://doi.org/10.1177/1745691610393525

Rottman, B. M., & Keil, F. C. (2012). Causal structure learning over time: Observations and interventions. *Cognitive Psychology*, *64*(1-2), 93-125. https://doi.org/10.1016/j.cogpsych.2011.10.003

Salmon, W. C. (1984). *Scientific Explanation and the Causal Structure of the World*. Princeton: Princeton University Press.

Salmon, W. C. (1994). Causality without counterfactuals. *Philosophy of Science*, *61*(2), 297-312.

Schlottmann, A. (1999). Seeing it happen and knowing how it works: How children understand the relation between perceptual causality and underlying mechanism. *Developmental Psychology, 35*(1), 303–317. https://doi.org/10.1037/0012-1649.35.1.303

Schlottmann, A., & Shanks, D. R. (1992). Evidence for a distinction between judged and perceived causality. *The Quarterly Journal of Experimental Psychology*, *44*(2), 321-342. https://doi.org/10.1080/02724989243000055

Scholl, B. J., & Nakayama, K. (2004). Illusory causal crescents: Misperceived spatial relations due to perceived causality. *Perception*, *33*(4), 455-469. https://doi.org/10.1068/p5172

Scholl, B. J., & Tremoulet, P. D. (2000). Perceptual causality and animacy. *Trends in Cognitive Sciences*, *4*(8), 299-309. https://doi.org/10.1016/S1364-6613(00)01506-0

Sloman, S. A., & Lagnado, D. A. (2005). Do we "do"? *Cognitive Science, 29*, 5–39.

Sloman, S. A., & Lagnado, D. (2015). Causality in thought. *Annual Review of Psychology*, *66*(1), 223-247. https://doi.org/10.1146/annurev-psych-010814-015135

Sobel, D. M., & Kirkham, N. Z. (2006). Blickets and babies: The development of causal reasoning in toddlers and infants. *Developmental Psychology, 42*(6), 1103–1115. https://doi.org/10.1037/0012-1649.42.6.1103

Soo, K. W., & Rottman, B. M. (2018). Causal strength induction from time series data. *Journal of Experimental Psychology: General*, *147*(4), 485-513. https://doi.org/10.1037/xge0000423

Steyvers, M., Tenenbaum, J. B., Wagenmakers, E. J., & Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science*, *27*(3), 453-489. https://doi.org/10.1207/s15516709cog2703_6

Strickland, B., Silver, I., & Keil, F. C. (2017). The texture of causal construals: Domain-specific biases shape causal inferences from discourse. *Memory & Cognition*, *45*(3), 442-455. https://doi.org/10.3758/s13421-016-0668-x

Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind: Statistics, structure, and abstraction. *Science*, *331*(6022), 1279-1285. https://doi.org/10.1126/science.1192788

Wagemans, J., Van Lier, R., & Scholl, B. J. (2006). Introduction to Michotte's heritage in perception and cognition research. *Acta Psychologica*, *123*(1-2), 1-19. https://doi.org/10.1016/j.actpsy.2006.06.003

Waismeyer, A., & Meltzoff, A. N. (2017). Learning to make things happen: Infants' observational learning of social and physical causal events. *Journal of Experimental Child Psychology*, *162*, 58-71. https://doi.org/10.1016/j.jecp.2017.04.018

Waldmann, M. R., & Hagmayer, Y. (2005). Seeing Versus Doing: Two Modes of Accessing Causal Knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *31*(2), 216–227. https://doi.org/10.1037/0278-7393.31.2.216

Waldmann, M. R., & Mayrhofer, R. (2016). Hybrid causal representations. In *Psychology of Learning and Motivation* (Vol. 65, pp. 85-127). Academic Press.

White, P. A. (2014). Singular clues to causality and their use in human causal judgment. *Cognitive Science*, *38*(1), 38-75. https://doi.org/10.1111/cogs.12075

White, P. A., & Milne, A. (1997). Phenomenal causality: Impressions of pulling in the visual perception of objects in motion. *The American Journal of Psychology*, *110*(4), 573-602. https://doi.org/10.2307/1423411.

Wolff, P. (2014). Causal pluralism and force dynamics. In B. Copley, & F. Martin (Eds.), *Causation in Grammatical Structures* (pp. 100-119). Oxford, UK: Oxford University Press.

Woodward, J. (2005). *Making Things Happen: A Theory of Causal Explanation*. Oxford, UK: Oxford University Press.

Woodward, J. (2011). Causal perception and causal cognition. In J. Roessler, H. Lerman, & N. Eilan (Eds.), *Perception, Causation, and Objectivity* (pp. 229-263). Oxford, UK: Oxford University Press.

Yela, M. (1952). Phenomenal causation at a distance. *Quarterly Journal of Experimental Psychology*, *4*(4), 139-154. https://doi.org/10.1080/17470215208416612

Yeung, S., & Griffiths, T. L. (2015). Identifying expectations about the strength of causal relationships. *Cognitive Psychology*, *76*, 1-29. https://doi.org/10.1016/j.cogpsych.2014.11.001