



Cognitive Science 44 (2020) e12871

© 2020 The Authors. *Cognitive Science* published by Wiley Periodicals LLC on behalf of Cognitive Science Society (CSS). All rights reserved.

ISSN: 1551-6709 online

DOI: 10.1111/cogs.12871

Time and Singular Causation—A Computational Model

Simon Stephan,  Ralf Mayrhofer, Michael R. Waldmann 

Department of Psychology, University of Göttingen

Received 10 January 2020; received in revised form 19 May 2020; accepted 27 May 2020

Abstract

Causal queries about singular cases, which inquire whether specific events were causally connected, are prevalent in daily life and important in professional disciplines such as the law, medicine, or engineering. Because causal links cannot be directly observed, singular causation judgments require an assessment of whether a co-occurrence of two events c and e was causal or simply coincidental. How can this decision be made? Building on previous work by Cheng and Novick (2005) and Stephan and Waldmann (2018), we propose a computational model that combines information about the causal strengths of the potential causes with information about their temporal relations to derive answers to singular causation queries. The relative causal strengths of the potential cause factors are relevant because weak causes are more likely to fail to generate effects than strong causes. But even a strong cause factor does not necessarily need to be causal in a singular case because it could have been preempted by an alternative cause. We here show how information about causal strength and about two different temporal parameters, the potential causes' onset times and their causal latencies, can be formalized and integrated into a computational account of singular causation. Four experiments are presented in which we tested the validity of the model. The results showed that people integrate the different types of information as predicted by the new model.

Keywords: Singular causation; Causal attribution; Preemption; Time; Causal reasoning; Computational modeling

Correspondence should be sent to Simon Stephan, Department of Psychology, University of Göttingen, Gosslerstr. 34, 37073 Göttingen, Germany. E-mail: simon.stephan@psych.uni-goettingen.de

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

1. Introduction

The ability to reason causally about how the world works is one of our central cognitive capacities (see Sloman, 2005; Waldmann, 2017, for overviews). While previous research has frequently investigated the question of how we learn and reason about general causation, that is, how we learn about causal relationships referring to classes of objects or people (e.g., “smoking causes lung disease”), relatively little is known about how we reason about causal relations between singular events (but see Danks, 2017; Lombrozo & Vasilyeva, 2017). For example, if a smoker is diagnosed with lung cancer, we might wonder whether it has been her smoking that caused her lung cancer or whether the disease was caused by something else, such as exposure to asbestos. In the context of law, we may be interested in whether a defendant actually caused the death of his fiancée when he was drunk, which is a different query from a sociologist who is interested in the general relation between alcoholism and violent acts (see Hart & Honoré, 1959/1985; Lagnado & Gerstenberg, 2017). Singular causation queries are aided by knowledge about general causal relations, but cannot be reduced to them. Knowing that smoking probabilistically causes lung cancer does not entail that the lung cancer of a specific smoker was actually caused by her smoking.

More abstractly, when assessing singular causation, reasoners need to take into consideration that a singular co-occurrence c and e of the general factors C and E , even if it is known that there is a general causal relation between C and E (i.e., $C \rightarrow E$), could well be a mere coincidence. The present research asks how reasoners form beliefs about causal links in singular cases. If the potential causes and a particular outcome are known to have co-occurred, what influences the degree to which a reasoner believes that there has been a causal connection between them?

Previous studies on singular or actual causation have focused mostly on a different problem, causal selection, which refers to the phenomenon that even though several factors typically need to come together to generate an effect, people have the strong tendency to regard only a subset of these factors as the cause(s) of the observed outcome (e.g., Cheng & Novick, 1991; Hitchcock & Knobe, 2009; Icard, Kominsky, & Knobe, 2017; Kominsky, Phillips, Gerstenberg, Lagnado, & Knobe, 2015; Lagnado, Gerstenberg, & Zultan, 2013; Novick & Cheng, 2004; Phillips, Luguri, & Knobe, 2015). An often-discussed example is a forest fire that occurred after a lit match had been dropped (e.g., Halpern & Hitchcock, 2015). In this example, most people are more inclined to regard the dropping of the lit match as the actual cause of the fire than the presence of oxygen in the surrounding atmosphere. Importantly, selecting the lit match over oxygen presupposes that the lit match has already been identified as being causally linked to the effect in this particular situation. In this article, we do not address the question of what the factors are that determine whether an event that is already known to be causally linked to the effect is selected over others. We are interested in the upstream question of how the existence of a singular causal link can be established in the first place.

To explain how reasoners determine whether two co-occurred events c and e were causally connected and did not co-occur merely coincidentally, we propose a new

computational model of singular causation judgments that is based on the power PC framework of causal attribution developed by Cheng and Novick (2005) and extended by Stephan and Waldmann (2018). We will show that the assessment of singular causation relations requires the combination of two types of information: (a) information about the (general) causal strengths of the potential cause factors, and (b) information about temporal features of the potential causes and about their temporal relation. The model proposed by Stephan and Waldmann (2018) left open the questions of which temporal factors are relevant for singular causation, how these factors can be assessed, and how they can be computationally integrated with causal strength information to derive predictions for singular causation judgments. The present paper addresses these open questions. We will focus on two types of temporal information: (a) information about the onset times of the potential causes, and (b) information about causal latencies, that is, the time it takes causes to bring about their effects. We will show how information about these components can be formalized and how the generalized model of causal attribution (Stephan & Waldmann, 2018) combines causal strength and temporal information to predict singular causation judgments. The results of four experiments will be reported that support the predictions of the new model. All experimental materials, including demo videos for the studies, and the behavioral data are available under <https://osf.io/5yvs4/>.

2. Assessments of singular causation require temporal information

As an everyday life example illustrating that assessments of singular causation require the consideration of both causal strength and temporal information, consider the question of whether Mary's current headache is a side effect of the particular medicine she had taken earlier that day or whether her headache was actually caused by some alternative cause. To answer this singular causation query, it is necessary to know whether this type of medicine generally tends to cause headaches as a side effect and how strong the general causal link is: The stronger it is, the more likely it is that a singular instantiation of the respective event types is causal rather than coincidental. Merely knowing the general causal strength of this relation is not enough, however, to determine whether Mary's having taken the medicine actually caused her headache, because even a perfectly reliable cause with deterministic causal strength could have been preempted in its efficacy by an alternative cause on a given singular occasion. Maybe Mary's headache was actually caused by the unusual yoga exercise she also did earlier on that day. In the present article, we argue that temporal information is relevant for the assessment of singular causation because assessments of singular causation must take the possibility of causal preemption into account—and it is through the combination of causal strength and temporal information that this goal can be accomplished.

Cases of preemption are an often-discussed phenomenon in the philosophical literature on singular causation (e.g., Halpern, 2016; Halpern & Hitchcock, 2015; Hitchcock, 2007, 2009; Lewis, 1973; Menzies, 1996; Pearl, 2000) because they pose a problem for

counterfactual and probabilistic theories of causality. A prominent example of preemption that can be used to illustrate the temporal dimensions that we consider to be relevant is a scenario involving two rock throwers, Billy and Suzy. It is assumed in this scenario that neither Billy nor Suzy, when acting alone, ever fails to hit and destroy a bottle. Hence, it is assumed that the rock throwings of either protagonist have deterministic general causal strength. It is then mentioned that, on a particular occasion, Billy and Suzy both end up aiming for the same bottle and both are throwing their rocks with nearly identical speed. Crucially, Suzy manages to throw her rock a bit earlier than Billy, and we hear the bottle shatter. In this case, it was intuitively Suzy's and not Billy's rock throwing that was the singular cause of the bottle's breaking, regardless of our certainty that Billy's throw would have broken the bottle had Suzy not thrown her rock first.

2.1. Onset differences

One temporal factor relevant to determine the preemptive relation between the potential singular causes of an effect, which is made explicit in the scenario description about Billy and Suzy, is the *onset difference* between the potential causes. It is explicitly mentioned in the story that Suzy throws her rock earlier than Billy.

Generally, singular cases involving preemption can be illustrated schematically with neuron diagrams like the one shown in Fig. 1 (see Paul & Hall, 2013). The diagram models a situation like the one in the discussed scenario in which one of the potential causes, say our potential target cause c , was preempted by an alternative cause a that occurred earlier than c . Scenarios of causal preemption in which one of the potential causes is instantiated earlier than the other are also called “early preemption” in the philosophical literature (see, e.g., Hitchcock, 2007). Abstractly, what the scenario illustrates is that, everything else being equal, if a occurs earlier than c on an occasion on which both are sufficiently strong to generate the effect, then a will causally preempt c . Conversely, c will causally preempt a when c occurs earlier than a in this kind of situation. These two possibilities are depicted in Fig. 2A.

2.2. Causal latency

We have mentioned above that onset differences between the potential singular causes of an effect determine their preemptive relation under the assumption that “everything




Fig. 1. Neuron diagram modeling a situation of causal preemption (cf. Paul & Hall, 2013). The interrupted causal arrow departing from c and the arrow connecting a and e illustrate a situation in which c was causally preempted by a . The fact that a is positioned left of c and the fact that c is positioned left of e indicates that a occurred first, followed by c and then e .

else is equal.” The second temporal factor that we suggest to be relevant for the assessment of singular causation and to which the “everything else being equal” clause is referring is the potential causes’ *causal latencies*, by which we mean the time it takes a cause to produce its effect. For example, swallowing a dose of aspirin does not bring relief right away. It takes some time until the medicine manifests its pain-relieving capacity. When we enter an elevator and press a button, it takes a moment until the elevator begins to move. In the scenario about Billy and Suzy, information about causal latency was conveyed when it was mentioned that both protagonists are throwing their rocks toward the bottle with almost identical speed.

An important characteristic of causal latencies is that they are, like causal strength and unlike event onsets, typically unobservable properties of causes. The causal latency of a general cause factor C must be inferred from observable onset differences between singular instantiations of C and E (cf. Bramley, Gerstenberg, Mayrhofer, & Lagnado, 2018). In an environment that is shielded from potential alternative causes, the observed delay between c and e represents a direct measure of the causal latency of c ; otherwise C ’s causal latency has to be extracted from the background rate of E produced by alternative causes A .

Neuron diagrams illustrating the relevance of causal latency for the assessment of singular causation are depicted in Fig. 2B. Everything else being equal (i.e., given identical onset times of the potential causes) a sufficiently strong potential target cause c of the effect e preempts a simultaneously sufficiently strong potential alternative cause a on an occasion if c ’s causal latency is smaller than a ’s. Causal latency here refers to the time it would have taken the target cause c to produce the effect if it had acted in isolation. Conversely, a preempts c in causing e if a ’s causal latency is smaller than c ’s on this occasion. Of course, the “everything else being equal” criterion will rarely hold in real-life situations. On most occasions, different onset times will be combined with different causal latencies, which leads to the possibility that even causes with a long causal latency can preempt alternative causes if they have an onset advantage. Likewise, a target cause c that occurs after its causal competitors can still manage to preempt them if it has a short causal latency.




Fig. 2. Illustration of the temporal components relevant for the assessment of the potential preemptive relation between two potential causes of an effect. (A) illustrates the relevance of the instantiation times of the potential causes. (B) illustrates the relevance of the causal latencies of the potential causes.

The focus of the present paper is on the role of temporal assumptions in singular causation judgments. However, given that singular causation assessments require general causation knowledge (cf. Danks, 2017), it is interesting to see what role temporal assumptions have been found to play in general causal learning (see Buehner, 2017, for a state-of-the-art review of research on the role of time in general causation learning). A typical early finding was that reasoners seem to have a strong preference for spatio-temporal contiguity when learning about the general strengths of causes (e.g., Shanks, Pearson, & Dickinson, 1989). These studies manipulated temporal gaps and found that learning about contingencies becomes more difficult with increasing spatio-temporal distance between causes and effects.

However, it could also be shown that contiguity becomes less potent once background knowledge about latencies is activated by means of cover stories drawing on prior knowledge. For example, Hagmayer and Waldmann (2002) showed that assumptions about temporal delays determine which events in a sequence of observed events are considered candidates for a causal relation, which in turn affected which contingencies were learned as indicators of causal strength.

Another study, conducted by Buehner and McGregor (2006), confronted subjects with an unfamiliar apparatus in which marbles running down a pathway could activate different switches connected to a light bulb. During an inspection phase, the steepness of the pathways within the machine was varied, which either suggested a relatively short or a relatively long causal latency. During the test phase, the mechanism of the device was covered and subjects observed multiple instances of marble insertion and light activation. The observed delays were either compatible or incompatible with previously learned causal mechanisms. The results showed that subjects who observed highly contiguous successions of events only gave high causality ratings when the mechanism in fact suggested short delays, while causality ratings were low when subjects expected the delays to be longer.

In another study, Lagnado and Speekenbrink (2010) showed that lowered causal strength judgments with increasing causal delays might indeed be due to subjects' assumptions about the preemptive relation between the target cause and potential alternative causes that occur in the delay period. In their experiment, Lagnado and Speekenbrink (2010) varied the causal latency of the target cause factor and the probability with which alternative causes of the effect occurred during the cause–effect interval. This experimental design made it possible to compare causal judgments between conditions with equal causal latencies but varying probabilities with which potential alternative causes occurred prior to the effect. Causal judgments tended to be lower when the probability of the occurrence of alternative causes prior to the occurrence of the effect was high but not when this probability was low, presumably because the occurrence of alternative causes before the occurrence of the effect opens up the possibility that the target cause was preempted. Variation in causal latency alone did not affect causal judgments in this study. While the study of Lagnado and Speekenbrink (2010) focused on the role of temporal relations and possible cases of preemption in a general causation learning task, our focus

is on the interaction between causal strength, temporal assumptions, and preemption in judgments about singular causation relations.

3. Combining causal strength and temporal factors—The generalized power model of causal attribution

The crucial idea of the present paper is that the assessment of singular causation requires the integration of knowledge about general causal strength of the potential causes of an observed effect and about the temporal relation between them. We have already introduced the temporal factors that we consider to be crucial, the onset difference between the potential causes and their causal latencies. Building on the power PC framework of causal attribution (Cheng & Novick, 2005; Stephan & Waldmann, 2018), we will show how all these factors can be combined in an integrated model of singular causation judgments. As in the scenarios discussed in the introduction, our analysis focuses on the relatively simple case in which a target effect E can be produced by two potential causes, either by a target cause C or by an alternative cause A . Moreover, we consider singular occasions on which it is known that all factors were actually instantiated, that is, on situations in which it is observed that $C = 1$ (or c in short), that $A = 1$ (i.e., a), and that $E = 1$ (i.e., e).

The general causal relation between C , A , and E can be modeled with the common-effect causal Bayes net depicted in Fig. 3A, in which b_C and b_A denote the base rates of the two causes and w_C and w_A denote the causal strength of C and A , respectively. It is assumed that C and A represent independent causes of E that combine their influences according to a noisy OR-gate (Griffiths & Tenenbaum, 2005; Pearl, 1988).¹ According to the causal Bayes net framework (Pearl, 1988, 2000; Spirtes, Glymour, & Scheines, 1993) and rational models of causal induction (Cheng, 1997; Griffiths & Tenenbaum, 2005), causal strength can be understood as the probability with which a cause generates the effect E independent of the influence of the other causes of E (see also Rottman, 2017; Rottman & Hastie, 2014, for overviews of research on causal Bayes nets as models of causal reasoning).

Based on Cheng's (1997) causal power PC theory, Cheng and Novick (2005) have proposed a computational model which we will call the *standard* power model of causal attribution here. The model shows how knowledge about the general causal strengths of the potential causes of an observed singular effect e can be used to determine whether a potential target cause c was the singular cause of e . According to the standard model, the probability that c caused e on a singular occasion on which c , e , and a potential alternative cause a are known to have occurred (cf. Fig. 3B) can be denoted $P(c \rightarrow e|c, a, e)$, which can be computed using Eq. 1:

$$P(c \rightarrow e|c, a, e) = \frac{w_C}{w_C + w_A - w_C \cdot w_A} = \frac{w_C}{P(e|c, a)}. \quad (1)$$

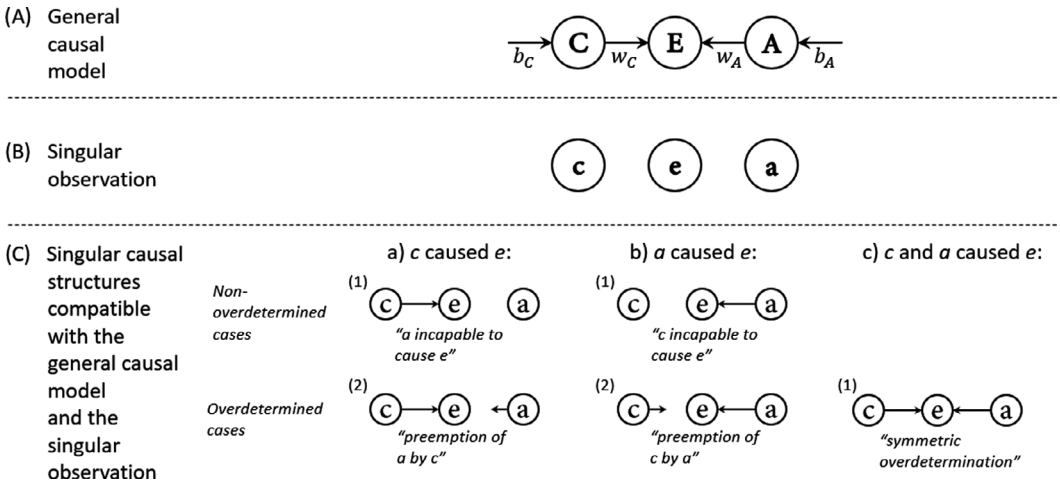


Fig. 3. Illustration of a general common-effect model (A) with two independent causes C and A , underlying a situation in which c , a , and e were observed (B). (C) shows alternative singular causal structures compatible with the general causal structure and the singular observation.

In this equation, the target cause’s causal strength, w_C , is normalized by the probability of the effect in the presence of its potential causes, $P(e|c, a, e)$. $P(e|c, a, e)$ corresponds to the sum of the potential causes’ causal strengths minus their intersection. Stephan and Waldmann (2018) have argued that the model effectively represents a “sufficiency filter”: It identifies among all co-occurrences of c , a , and e the proportion of cases in which the target cause was sufficiently strong to generate the effect. The standard model thus correctly captures the notion that the probability that a target cause is the singular cause of an effect increases with the target cause’s general causal strength. However, as Stephan and Waldmann (2018) have pointed out, a consequence of the model’s focus on causal strength is that it neglects the possibility that the target cause could have been preempted by the alternative cause on a particular occasion. As a result, Eq. 1 tends to overestimate $P(c \rightarrow e|c, a, e)$. For example, for the rock-throwing scenario in which Billy and Suzy are assumed to be perfectly accurate rock throwers but Suzy acts earlier, the model makes the counterintuitive prediction that we should be certain that both Suzy’s and Billy’s throwing were the singular cause of the effect. The scenario can be modeled by setting $w_C = w_A = 1$, which would yield $P(c \rightarrow e|c, e, a) = \frac{1}{1} = 1$.

Figure 3C shows the different singular causal relations that are compatible with the assumed general common-effect model shown in Fig. 3A and the singular observation of the three events c , a , and e , shown in Fig. 3B. Since the standard model solely focuses on the causal strengths of the potential causes, it correctly attributes causality to c in the situations $a1$, $a2$, and $c1$. Situation $a1$ illustrates a scenario in which the alternative cause a does not possess sufficient power to generate the effect. In Scenario $a2$, the target cause c preempts its competitor a . In Scenario $c1$, both c and a act in a way that is called “symmetric overdetermination” in the literature (cf. Hitchcock, 2001; Paul & Hall, 2013;

Pearl, 2000; Strevens, 2008); that is, both are sufficiently strong to generate the effect and they unfold their causal capacity synchronously. However, the model also erroneously attributes causality to c whenever a situation exhibits structure b_2 , in which a preempts c in its efficacy.

To incorporate the possibility of causal preemption, Stephan and Waldmann (2018) have proposed a generalization of the standard power PC model of causal attribution. Eq. 2 summarizes the generalized model:

$$P(c \rightarrow e|c, a, e) = \frac{w_C - w_C \cdot w_A \cdot \alpha}{w_C + w_A - w_C \cdot w_A} = \frac{w_C \cdot (1 - w_A \cdot \alpha)}{P(e|c, a)}. \quad (2)$$

Equation 2 introduces a new parameter, α , which is weighted with the product of w_C and w_A . The product $w_C \cdot w_A \cdot \alpha$ captures the probability of preemption of the target cause by its competing alternative cause. The intersection of w_C and w_A identifies cases in which C and A are both probabilistically sufficient for E . This part of the term is relevant because causal preemption can only occur when the potential causes are each probabilistically powerful enough to generate the effect. On these occasions, it can either be the case that c preempts a , that a preempts c , or that c and a act symmetrically and instantiate a case of “symmetric overdetermination.” The α parameter represents an allocation parameter, $0 \leq \alpha \leq 1$, that determines the proportion of the sufficiency overlap in which C is preempted by A . As $w_C \cdot w_A \cdot \alpha$ captures the probability of preemption of the target cause by its competitor, it needs to be subtracted from w_C .

To demonstrate deviations of subjects’ singular causation judgments from the predictions made by the standard model, Stephan and Waldmann (2018) used a cover story in their study in which it was plausible that the alternative causes fully preempted the target cause in its efficacy (captured by the new model by an α value of 1). The experiments supported the model. However, a shortcoming of the study was that the temporal factors underlying a possible case of preemption were neither spelled out nor experimentally manipulated. The present article remedies this shortcoming by providing a link between the value of the α parameter and the temporal factors onset differences and causal latencies.

3.1. Quantifying alpha based on onset differences and causal latencies

The onset times of the potential causes c and a are in the following abbreviated as t_c and t_a . The onset difference between them can be denoted as $\Delta_t = t_a - t_c$. In Fig. 2A, Δ_t is shown on the x -axes. Formally, if two potential causes c and a are both sufficiently powerful on an occasion and have identical causal latencies, c preempts a in causing e when Δ_t is positive. Conversely, a preempts c in causing e when Δ_t is negative.

As for causal latencies, in some contexts we may experience causal latencies that have an (almost) invariant value (e.g., the causal delay between the pressing of an elevator button and the elevator’s starting to move), but in most cases causal latencies will be variable. For example, a dose of aspirin might normally take effect after 20 min, but it is

reasonable to expect some variability (explained by the operation of unobserved hasteners or delayers; see, e.g., Lagnado & Speekenbrink, 2010). In most cases, it therefore seems appropriate to represent causal latency as a random variable following a particular distribution.

A standard way of modeling variable latencies is to use gamma distributions. Gamma distributions are used, for example, in queueing theory to model waiting times (Shortle, Thompson, Gross, & Harris, 2018). In a recent study in which Bramley et al. (2018) investigated the role of time in general causal structure induction, gamma distributions were used to model causal latencies. The gamma distribution represents a generalization of the exponential distribution and is characterized by two parameters: shape, $\kappa > 0$, and scale, $\theta > 0$. The expected value of a random variable X following a gamma distribution is $E[X] = \kappa \cdot \theta$. Its variance is $Var[X] = \kappa \cdot \theta^2$. Different gamma distributions together with their expected values (vertical lines) are shown in Fig. 4.

With these formalizations of the different temporal components at hand, the α parameter of our model can be quantified. We here consider situations in which the causal latencies of the potential causes have been learned through multiple observations and in which the causes' onset difference in the target situation is known. In these cases, the probability that the target cause c was preempted by its competitor a if both happen to be simultaneously sufficient for the effect corresponds to:

$$\alpha = P(t_{a \rightarrow e} + \Delta_t < t_{c \rightarrow e} | e, c, a). \quad (3)$$




Fig. 4. Example of gamma distributions with different shape (κ) and scale (θ) parameters. The expected value of a gamma distribution is given by $E[X] = \kappa \cdot \theta$. The expected values of the example distributions are illustrated by the vertical lines.

α corresponds to the probability that the sum of the causal latency of the competing cause a and the difference between c 's and a 's onset times is smaller than the causal latency of c , given e , c , and a . If the causal latencies of C and A are represented as fixed invariant values, α can be determined simply by comparing the sum of $t_{(a \rightarrow e)}$ and Δ_t with $t_{(c \rightarrow e)}$. In this case, α will either be 0 or 1. If the causal latencies exhibit variability and are modeled with gamma distributions, α can take on any value between 0 and 1. Its value can then be approximated with the following Monte Carlo (MC) algorithm:

1. Sample N pairs of causal latencies ($t_{c \rightarrow e}$, $t_{a \rightarrow e}$) from the latency distributions ($t_{C \rightarrow E}$, $t_{A \rightarrow E}$) of C and A , respectively.
2. Calculate $t_{a \rightarrow e} + \Delta_t$ for all sampled $t_{a \rightarrow e}$ -values.
3. Count all pairs for which $t_{a \rightarrow e} + \Delta_t < t_{c \rightarrow e}$.
4. Divide this count by N .

For an illustration, consider Fig. 4 again. Assume that the target cause's causal latency follows the gamma distribution with the parameters $\kappa = 10$ and $\theta = 150$, and that the alternative cause's causal latency corresponds to the rightmost distribution. It can be seen that the target cause tends to act more quickly than its competitor but that the curves also visibly overlap, which implies that there exist occasions on which the alternative cause outperforms the target cause. The α parameter captures the probability with which this happens. For occasions on which the two causes occur simultaneously (i.e., $\Delta_t = 0$), the MC algorithm presented above yields a value of $\alpha = 0.05$, which means that the probability that a acts more quickly than c on such a singular occasion is 5%. Now imagine that it is known that both causes not only occurred at the same time but also that their causal latencies followed the same gamma distribution. In this case, α would be 0.50 because in a set of randomly sampled pairs of causal latencies, the sampled latency of the alternative cause would be smaller in half of the pairs.²

Let us now demonstrate how the whole generalized model works by first looking again at the original version of the rock-throwing scenario in which the protagonists represent deterministic causes. The model's causal strength parameters w_C and w_A need to be set to 1.0 in this case. Furthermore, since Billy and Suzy both throw their rocks with almost identical speed on the target occasion, identical causal latency values can be assumed in this case, $t_{c \rightarrow e} = t_{a \rightarrow e}$. Finally, since Suzy throws her rock earlier than Billy, Δ_t is < 0 . α would take on a value of 1.0 in this case because $t_{a \rightarrow e} + \Delta_t < t_{c \rightarrow e}$. If we insert all parameter values into Eq. 2, we obtain $P(c \rightarrow e | c, a, e) = 0$. Hence, in line with the intuition the scenario aims to convey and contrary to what the standard model predicts, the generalized model predicts that Billy's rock throwing was not the singular cause of the bottle's shattering.

Let us now consider the more complicated probabilistic case and assume that rock throwings of Billy and Suzy each have a causal strength of only 0.5, that is, $w_C = w_A = 0.50$ this time. We further assume that the velocities with which the protagonists throw their rocks exhibit variability and that Billy, on average, throws his rocks slightly faster than Suzy. We assume that Billy's and Suzy's causal latencies are described by the middle and the rightmost gamma distribution in Fig. 4, respectively. Finally, we assume that

Suzy throws her rock exactly 550 ms before Billy on the target occasion, which yields $\Delta t = -550$ ms. Remember that when we first considered the two rightmost gamma distributions in Fig. 4 and assumed $\Delta t = 0$ ms, we obtained $\alpha = 0.05$. Since we now assume that $\Delta t = -550$ ms, we can imagine that the rightmost gamma distribution describing Suzy's causal latency is shifted 550 ms to the left on the x-axis, thereby increasing the overlap of the distributions. As a result, it should be more likely now that Suzy's throwing her rock preempted Billy's. Applying our sampling algorithm, we indeed obtain a higher value of 0.20 for α .

We can now illustrate what happens if we enter all the parameter values into Eq. 2. First, $w_C \cdot w_A$ in the numerator yields $0.50 \cdot 0.50 = 0.25$, implying that Billy and Suzy's rock throwings are simultaneously sufficiently powerful to generate the effect in 25% of the situations in which the two agents aim at the same target. The α value of 0.20 that we estimated implies that Suzy's throwing can be expected to preempt Billy's in 20% of *these* situations. We thus obtain $w_C \cdot w_A \cdot \alpha = 0.50 \cdot 0.50 \cdot 0.20 = 0.05$ for the product in the equation's numerator. The product of 0.05 means that Billy's throwing is preempted by Suzy's throwing in 5% of all situations in which both protagonists aim at the same bottle. These 5% should therefore be attributed away from the target cause according to the model. We obtain $P(c \rightarrow elc, a, e) = 0.60$: The probability that Billy's throwing was the singular cause of the effect on this occasion is 60%.

Let us now see what would have happened if Suzy had not had the onset advantage on her side and that both protagonists had thrown their rocks simultaneously. We already know that α is 0.05 in this case, implying that we should have been slightly more confident that Billy's throwing was the singular cause. The model predicts that our confidence should have increased to $P(c \rightarrow elc, a, e) = 0.65$. Now imagine it had been Billy who had the onset advantage on his side, implying $\Delta t = 550$ ms. α would have been close to zero in this case (0.01). Yet $P(c \rightarrow elc, a, e)$ would still have only increased to 0.66. The reason for this rather small increase is the relatively low causal strength of Billy's throwing of only 0.50 and the fact that the rock throwing of his competitor Suzy has identical causal strength. The standard model (Eq. 1) predicts almost the same value in this case (0.67), which illustrates that the generalized model reduces to the standard model when α approaches 0.

Consider now the case of "symmetric overdetermination." Imagine that Billy and Suzy throw their rocks with identical speed at the same time with sufficient strength. This case can be modeled by setting α to 0. If both potential causes are by definition assumed to act symmetrically, no preemption occurs and it seems normative to judge both acts to be singularly causally linked to e (cf. Lagnado & Gerstenberg, 2017; Pearl, 2000; Schaffer, 2003).³

In sum, the generalized power model of causal attribution we have presented combines causal strength information with temporal information about onset times and causal latencies to estimate the probability that a target cause was the singular cause of an observed effect. Fig. 5 shows the predictions that the model makes for different value combinations of w_C , w_A , and α . Colored lines represent different values of the target cause's causal strength w_C ; the alternative cause's causal strength w_A is shown on the x-axes (decreasing

in magnitude). The different panels represent different α values. As can be seen, the model predicts that a reasoner should be more confident that the target cause was the singular cause of the observed effect the higher its causal strength is relative to the causal strength of the potential alternative cause. Given a particular non-zero value for w_A , $P(c \rightarrow elc, a, e)$ increases with increasing values of w_C . Similarly, given a particular non-zero value for w_C , $P(c \rightarrow elc, a, e)$ increases with decreasing values for w_A . In the case of $w_A = 0$ and $w_C > 0$, the model predicts that we can be certain that c caused e because there is no alternative cause that could have generated e . Fig. 5 also shows that these effects are moderated by the probability with which the target cause is preempted by the alternative cause: Given a particular combination of w_C and w_A , our confidence that the target cause caused the effect should decrease with increasing values of α . It can also be seen that this mitigating effect is particularly pronounced for high causal strength values. For example, for $\alpha = 1.0$, which is shown in the last panel, $P(c \rightarrow elc, a, e)$ approaches 0 the stronger the alternative cause A is. For $\alpha = 0$ (first panel), the predictions of the generalized model are identical to those made by the standard model. This panel with $\alpha = 0$ illustrates that the standard model, due to its neglecting temporal information and the possibility of preemption, tends to overestimate $P(c \rightarrow elc, a, e)$.

The following four experiments will test the core assumptions of our generalized model. The focus of the studies will be on the interaction between causal strength, temporal assumptions, and preemption in judgments about singular causation relations.

4. Experiment 1

The goal of Experiment 1 was to test whether the two temporal components, causal latency, and onset difference, affect reasoners' singular causation judgments in the way predicted by the new model. The α parameter of the new model integrates onset and causal latency information and assigns equally important roles to both factors, which means that a relatively slowly acting cause might compensate its disadvantage on this dimension with an advantage on the onset dimension and vice versa. We aimed to test if subjects would integrate both types of information in this way. Moreover, the first study aimed to isolate the influence of the temporal factors from the influence of causal strength, which was achieved by testing deterministic causes. The standard causal power model of attribution (Eq. 1) predicts high singular causation ratings in this case. Finally, to keep things simple in this initial study, the cause factors were assigned fixed, non-varying causal latencies as well as fixed, non-varying onset times.

The four different scenarios involving two potential causes, labeled A and B , of a single effect are shown in Table 1. The two causes were associated with different combinations of causal latency and onset times across the different situations. When both potential causes have equal causal strengths, our model predicts that subjects should give high singular causation ratings for the cause for which the sum of causal latency and onset time is smaller, as it is unlikely that this cause was preempted by its competitor. The sum of causal latency and onset time for each of the two potential causes is listed in the last row of Table 1. Subjects




Fig. 5. Predictions of the generalized power model of causal attribution (Eq. 2) for different value combinations of w_C , w_A , and α .

were expected to give high singular causation judgments when the target cause was Cause A and low judgments when the target cause was Cause B.

In the first scenario, the influence of onset differences was isolated as both causes had identical causal latencies. While both potential causes always produced the effect after 800 ms in this scenario, the onset difference was $\Delta_t = 800$ ms. The second scenario tested the influence of causal latencies. Both causes always occurred simultaneously 800 ms ($\Delta_t = 0$ ms) after a trial had begun but produced the effect with different latencies.

The last two scenarios (3 and 4) pitted causal latency and onset difference against each other. Two different versions (A and B) tested whether subjects actually integrated both temporal components or whether they considered one component to be more important than the other. For example, had we only tested the third scenario, higher singular causation judgments for cause A could be explained by a bias toward relatively low causal

Table 1
Overview of the set of different situations tested in Experiment 1

	(1) Different Onsets but Equal Causal Latencies		(2) Equal Onsets but Different Causal Latencies		(3) Different Onsets and Causal Latencies (A)		(4) Different Onsets and Causal Latencies (B)	
	Cause A	Cause B	Cause A	Cause B	Cause A	Cause B	Cause A	Cause B
Onset (ms)	800	1,600	800	800	1,600	800	800	2,400
Causal latency (ms)	800	800	800	1,600	800	2,400	1,600	800
Sum	1,600	2,400	1,600	2,400	2,400	3,200	2,400	3,200

latencies. Likewise, results in the expected direction in Scenario 4 could be explained by a bias toward causes with an onset advantage. By testing both cases, we can show that onset times and latencies are both relevant.

4.1. Methods

4.1.1. Participants

Three hundred and eighty-four subjects ($M_{\text{age}} = 35.14$, $SD_{\text{age}} = 12.38$, 202 female, 178 male, four subjects did neither indicate male nor female as their gender) were recruited via Prolific (www.prolific.ac). This sample size ($n = 96$ subjects per scenario) allows us to detect medium effects (of $d = 0.60$) in single group comparisons conducted with conservative two-sided t tests with more than 80% test power. We planned to analyze the data using planned contrasts, which entails an even higher test power. No data were analyzed until $n = 96$ subjects per condition had completed the study. Subjects were at least 18 years old, had at least an A-level degree, were native English speakers, and had an approval rate concerning participation in previous studies of at least 90%. Subjects were paid £1.25 for their successful participation.

4.1.2. Design, materials, and procedure

A demo video of the experiment is available under <https://osf.io/3rfgx/>. Subjects were randomly assigned to one of four different scenarios ($n = 96$) that varied with respect to the combination of causal latencies and onset differences of the two potential causes as shown in Table 1. Several factors were counterbalanced between subjects that will be described below.

Subjects were presented with a story about a fictitious medieval kingdom called “Extonia” whose king had two watchtowers (“North” and “South”) built at the border to protect the empire from invading barbarians. Subjects learned that the watchtowers were instructed to send carrier pigeons to the palace to cause alarm whenever they spotted barbarians approaching the border. It was pointed out that pigeons automatically set off an alarm the moment they arrive at the palace. Participants were asked to take the perspective of Extonia’s secretary of defense who plans to inspect the efficacy of the kingdom’s defense system. It was mentioned that two factors must be investigated: the flight durations of the pigeons from the two towers and the alertness of each tower (i.e., how fast each tower reacts when barbarians are coming). The flight durations of the pigeons constituted our manipulation of the towers’ causal latencies, whereas the reaction times of the towers allowed us to manipulate onset differences.

Participants were further instructed that they will learn about the two factors in separate learning phases. They were also informed that they will observe 10 pigeons from each tower in the flight duration phase (i.e., 20 observations in total) and that they will make 10 observations in the alertness phase. The function of these learning phases presenting multiple trials was to establish the general causal relationship and the relevant parameter values. The instructions additionally contained illustrations of how the

animations in the two different learning phases look (see Fig. 6). Finally, participants were informed that at the end they will be asked to answer a causal test query about a singular occasion on which the palace was alarmed. Subjects had to pass five instruction check questions probing their understanding of the scenario before they could proceed. Subjects who failed to answer all these questions correctly four times in a row were counted as invalid and excluded from the data set prior to any analysis of the data.

Whether subjects first learned the flight durations of the pigeons (cf. Fig. 6A) or began with the onset differences (cf. Fig. 6B) between the two towers was counterbalanced between subjects. The flight duration learning phase consisted of two parts, as subjects observed the two towers separately. Whether participants began with tower “North” or “South” was also counterbalanced between subjects. In each trial, the sending of a carrier pigeon was indicated by a colored circle surrounding the active watchtower that appeared with a delay of 500 ms. The corresponding inactive tower was transparent (see Fig. 6A). The arrival of the pigeon at the palace was signaled by a colored circle surrounding the palace. The circles remained visible until the end of the trial. The 10 flight durations presented for each tower were based on the causal latencies tested in the respective conditions. Whether tower “North” or tower “South” had the faster pigeons was counterbalanced between subjects. After each trial, participants had to click a “Next” button to proceed to the next observation. The “Next” button was operational 500 ms after the circle around the palace had appeared. After subjects finished their observations, they proceeded to the learning phase of the second relevant factor. The animation shown during the onset learning phase resembled the illustration depicted in Fig. 6B. The palace was masked with a grey rectangle in this phase to underline that it was irrelevant in this case. As in the flight duration phase, the sending of letter pigeons was indicated by colored circles surrounding the towers. The 10 onset differences subjects observed depended on condition. Whether tower “North” or tower “South” reacted more quickly was counterbalanced between subjects.

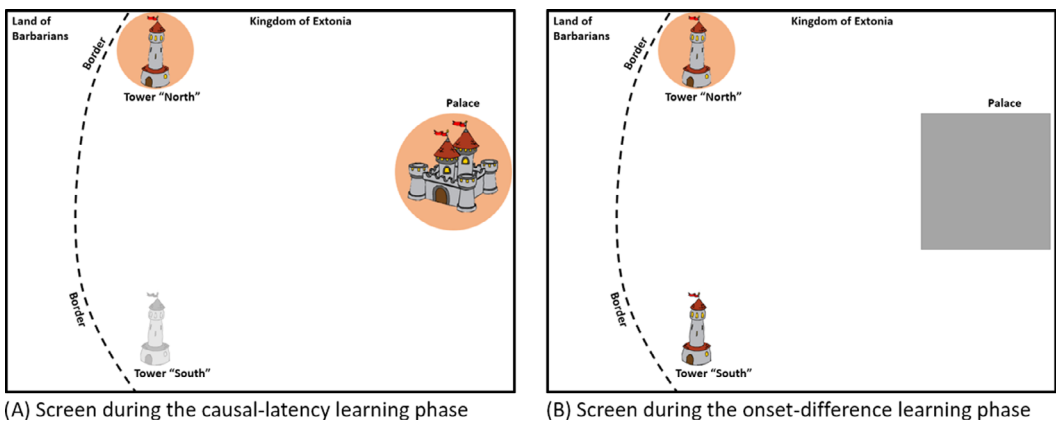


Fig. 6. Illustration of the learning task used in Experiment 1.

Subjects then proceeded to the target situation. A situation was described in which the palace had been alarmed. Participants read that the people of Extonia now want to decorate the crew of the tower that had caused the alarm on this occasion. All that the Extonians knew, however, was that both towers had spotted the barbarians and had sent a pigeon to the palace on this occasion. Subjects were also shown an image showing colored circles around each tower and around the palace. We then asked subjects to indicate how strongly they believed that the alarm was caused by tower “North”/“South” on this occasion. Judgments were provided using an 11-point rating scale with end points “definitely not caused by Tower ‘North/South’” and “definitely caused by Tower ‘North/South’” (the midpoint was labeled “50:50”). This rating was intended to measure $P(c \rightarrow \text{elc}, a, e)$.⁴ Whether the question referred to tower “North” or “South” was counterbalanced between subjects. A bipolar rating scale with one tower on either side of the scale would have forced subjects to decide for one of the towers, which would have introduced a disadvantage for the standard power PC model of causal attribution. As the standard model solely focuses on causal strength information, it predicts that both towers should be considered causal in this case. Finally, the orientation of the rating scale (i.e., whether high confidence had to be expressed on the left or on the right side of the scale) was counterbalanced between subjects.

4.2. Results and discussion

The results are summarized in Fig. 7. As can be seen, subjects made different singular causation judgments for the two potential causes in the four scenarios, which shows that they applied their background knowledge about both the onset differences and the causal latencies. The upper two panels in Fig. 7 show the results for the two scenarios that tested each temporal component in isolation (Scenarios 1 and 2 in Table 1). The left panel shows the results for Scenario 1 in which both causes had the same causal latency but different onsets, while the right panel shows the results for Scenario 2 in which both causes had identical onsets but different causal latencies. In the condition in which one cause factor tended to occur earlier than the other but both had identical causal latencies (Scenario 1) subjects who were asked about the “early” cause (Cause A in Table 1) gave higher singular causation judgments ($M = 0.75$, $SD = 0.24$, $Md = 0.80$, 95% CI [0.68, 0.82]) than subjects who were asked about the “late” cause ($M = 0.21$, $SD = 0.24$, $Md = 0.10$, 95% CI [0.14, 0.28]). A planned contrast revealed that this difference was significant, $t(376) = 9.9$, $p < .001$, $d = 2.25$. In the scenario in which both causes had synchronous onset times but varying causal latencies (Scenario 2), subjects gave higher singular causation ratings when asked about the cause with the shorter causal latency ($M = 0.83$, $SD = 0.21$, $Md = 0.90$, 95% CI [0.77, 0.89]) than when asked about the cause with the longer causal latency ($M = 0.25$, $SD = 0.25$, $Md = 0.10$, 95% CI [0.18, 0.32]). A planned contrast revealed that this difference was significant, $t(376) = 10.68$, $p < .001$, $d = 2.76$.

To test whether subjects actually integrated the two types of temporal information, ratings for Scenarios 3 and 4 were compared. We found that subjects tended to answer as

predicted by our model. In Scenario 3, subjects gave higher singular causation ratings when asked about the cause with the shorter causal latency ($M = 0.66, SD = 0.30, Md = 0.80, 95\% CI [0.57, 0.75]$) than when asked about the cause with the longer causal latency ($M = 0.47, SD = 0.32, Md = 0.50, 95\% CI [0.38, 0.56]$), even though the onset difference pointed in the other direction. A planned contrast testing this difference was significant, $t(376) = 3.56, p < .001, d = 0.63$. However, one possible explanation for this pattern is that subjects might have had a bias toward shorter causal latencies. The ratings obtained for Scenario 4 show that this does not explain the judgments, however, because subjects here gave higher singular causation ratings when asked about the cause with the longer rather than the one with the shorter causal latency ($M = 0.69, SD = 0.28, Md = 0.75, 95\% CI [0.61, 0.77]$ vs. $M = 0.42, SD = 0.25, Md = 0.50, 95\% CI [0.35, 0.49]$). A planned contrast testing this difference was significant, $t(376) = 4.87, p < .001, d = 0.96$. Subjects also had no general tendency to favor the cause that had an onset advantage, as one could have hypothesized based on the results of Scenario 4 alone. In Scenario 3, the




Fig. 7. Results (bars represent mean singular causation ratings; error bars denote 95% bootstrapped CIs) of Experiment 1.

cause with the onset advantage received lower ratings than the cause with the onset disadvantage. Additional evidence that subjects considered both temporal dimensions to be equally important is that the difference between the two causes is similar in Scenarios 3 and 4.

The results of this experiment show that subjects integrated onset and latency information when making singular causation judgments. Moreover, in line with what the generalized model predicts, subjects tended to favor the cause for which the sum of onset time and causal latency was minimal. The fact that the observed differences between the target causes were smaller in Scenarios 3 and 4 in which the causes differed on only one temporal dimension can be explained by the higher difficulty of the computation in the more complex cases and by the fact that the more complex conditions were more memory demanding. A further explanation, in line with Weber's law, is that the constant total difference of 800 ms between the respective competing causes (1,600 ms vs. 2,400 ms in Scenarios 1 and 2; 2,400 ms vs. 3,200 ms in Scenarios 3 and 4) might have been perceived as smaller in Scenarios 3 and 4.

5. Experiment 2

Experiment 1 shows that reasoners apply more than just their knowledge about the causal strength of the potential causes when assessing how likely it is that a particular factor C was causal on a singular occasion. Reasoners also rely on what they have learned about the onset times of the causes and the causes' causal latencies.

The causal latencies subjects learned in Experiment 1 had fixed non-varying values. The goal of Experiment 2 was to investigate and focus on the more natural case in which causal latencies are variable (while cancelling onset differences). We compared different conditions in which the two potential causes were paired with causal latencies described by different gamma distributions. On the assumption that variations of causal latencies are the default situation, we expected subjects to be sensitive to this factor. However, another possibility is that subjects might only learn and use the expected values of the causal latencies. The way in which we manipulated causal latencies in the present study allows us to test both possibilities (see Fig. 9). Again, we restricted the scenarios to deterministic causes and this time also dropped onset difference to be able to focus on the contribution of causal latency information.

The pairs of gamma distributions we contrasted in the five different conditions of the experiment are shown in Fig. 8. The distances between the distributions (G1–G5; higher numbers indicate higher expected values) in the different conditions varied quantitatively, from fairly extreme in Conditions 1 and 5 to identical in Condition 3. The latency distribution belonging to the target cause C in each condition is depicted in dark blue. In Condition 1, for example, the target cause's latency follows G1, whereas the latency of the alternative cause follows G5. Conditions 4 and 5 involve the same pairs of distributions as Conditions 1 and 2, but the distribution associated with the target cause was changed. Fig. 8 also shows the different α values estimated with the sampling algorithm presented

above. For the first pair, for example, α equals 0.01. In the case of deterministic causes, α corresponds to the probability that c was preempted by a . Thus, participants should be confident in this condition that it was the target cause c that caused the effect. In the fifth condition, by contrast, participants should be confident that c did not cause the effect. Fig. 8 also shows that all the other conditions should elicit more uncertainty. For example, in the third condition in which the potential causes have the same latency distribution, participants should be maximally uncertain about the singular cause of the outcome ($\alpha = 0.50$).

Figure 9 shows the predictions that the different models make. The predictions made by the standard model are shown in the first panel (Eq. 1). We also included the predictions derived from Eq. 2 with α set to 1, which are shown in the second panel. The goal here was to show that subjects do not misinterpret our test question as one that is asking for the probability that “ c alone” caused e on the target occasion, a possibility discussed by Cheng and Novick (2005). Note that the “ c alone” query according to Cheng and Novick’s (2005) framework asks for the probability with which c alone was probabilistically sufficient on an occasion because the standard model equates probabilistic sufficiency with causality. If subjects equated sufficiency with causality and understood the test query as a “ c alone” query, we should expect low ratings in all conditions (see second panel in Fig. 9). The third panel in Fig. 9 shows the predictions that the generalized model makes when α is estimated based on only the expected values of the causes’ causal latencies. In this case, the cause factor with the lower expected value of causal latency is

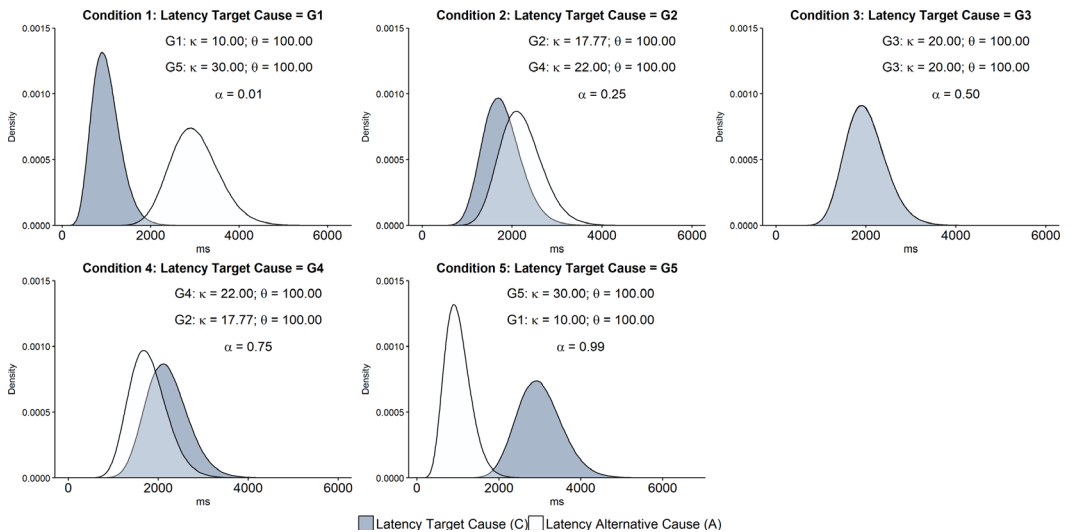


Fig. 8. Pairs of gamma distributions contrasted in the five conditions of Experiment 2. The shape (κ) and scale (θ) parameters of the five different gamma distributions are listed for each pair. The α values were obtained using the MC algorithm corresponding to Eq. 3. The dark distributions show the causal latencies of the target cause, and the light distributions show the causal latencies of the alternative cause.

identified as the singular cause. When both causes have the same expected causal latency value (Condition 3), the situation is treated as a case of symmetric overdetermination. Finally, the fourth panel shows the predictions the generalized model (Eq. 2) makes when α is estimated based on the respective causal latency distributions.

5.1. Methods

5.1.1. Participants

Two hundred subjects ($M_{\text{age}} = 27.14$, $SD_{\text{age}} = 8.71$, 119 female, 91 male) were recruited via Prolific (www.prolific.ac). This sample size ($n = 40$ subjects per condition) allows us to detect medium effects (of $d = 0.60$) in simple group comparisons using conservative t tests with more than 80% test power. The plan, however, was to analyze the data by fitting quantitative trends, which guarantees an even larger test power for medium effects. No data were analyzed until $n = 40$ subjects per condition had completed the study. The inclusion criteria were the same as in the previous study. Subjects were paid £0.80 for their successful participation.




Fig. 9. Model predictions (Plots 1–4) and results (Plot 5; bars show mean singular causation ratings; error bars denote 95% bootstrapped CIs) for the different conditions (see Fig. 8) of Experiment 2. The predictions of the standard model were obtained by setting alpha to 0. Plot 2 shows the predictions of the generalized model with α set to 1. Plots 3 and 4 show the generalized model’s predictions with different ways of estimating alpha: Plot 3 shows the predictions obtained when α is calculated based only on the expected values of the respective gamma distributions, and Plot 4 shows the predictions obtained when α is estimated based on the full gamma distributions.

5.1.2. Design, materials, and procedure

A demo video of the experiment is available under <https://osf.io/g6p72/>. Participants were randomly assigned to one of five conditions varying with respect to the contrasted gamma distributions and with respect to the gamma distribution associated with the target cause (see Fig. 8).

Because we included several counterbalancing factors, the full design had 40 conditions. The cover story and the paradigm were largely identical with the one in Experiment 1. This time, however, subjects were instructed that the flight durations tend to differ between different pigeons, and that they will therefore observe a sample of 13 pigeons from each tower. As in Experiment 1, before participants could proceed to the learning task, they had to pass five instruction check questions testing their understanding of the scenario.

During the causal latency learning task, the screen looked similar to the picture shown in Fig. 6A. The 13 flight durations presented for each tower corresponded to 13 quantiles that were uniformly distributed between the 5th and the 95th percentile of the respective gamma distribution (lists of the exact durations presented in this and in the other experiments can be accessed under <https://osf.io/k4mj9/>). This way, relatively small yet representative samples of the respective distributions could be presented. After each trial, participants had to click a “Next” button, which was operational 500 ms after the circle around the palace had been displayed. The 13 flight durations per tower were presented in random order. After subjects had observed both towers, they proceeded to an intermediate screen that informed them that they will next be shown the situation to which the test query refers.

Like in Experiment 1, the test scenario described a situation in which the palace was alarmed. Participants read that the people of Extonia wanted to decorate the crew of the tower that caused the alarm and that it was known that both towers had sent their pigeons simultaneously on this occasion. Subjects had to indicate how strongly they believed that the alarm was caused by tower “North”/“South”. Judgments were provided using an 11-point rating scale with the end points “definitely not caused by Tower ‘North/South’” and “definitely caused by Tower ‘North/South’” (the midpoint was labeled “50:50”). Whether the test question referred to tower “North” or tower “South” and the orientation of the scale were counterbalanced.

5.2. Results and discussion

The results are summarized in Plot 5 in Fig. 9 and in Table 2, which both show that singular causation ratings followed a negative linear trend. A polynomial trend analysis revealed that the observed negative linear trend was significant, $F(4, 195) = 111.70$, $p < .001$, $r = .83$. No other polynomial trend was significant.

The results are at odds with the predictions made by the standard model. As in Experiment 1, subjects used information about causal latency to make singular causation judgments. Moreover, subjects’ judgments closely followed the predictions of the version of our new model that utilizes gamma distributions to represent causal latencies (Eq. 2).

There was a high correlation between model predictions based on the gamma distributions and individual ratings, $r = .83$, $p < .001$, which is depicted graphically in the left scatterplot of Fig. 10. We also analyzed the fit between model predictions and mean singular causation judgments. The correlation between model predictions and group means is shown in the right panel of Fig. 10. This correlation amounts to $r = .99$, $p < .001$.

A limitation of the experiment is that the expected values of the gamma distributions and their variances were not manipulated independently; the gamma distributions with the higher expected values also had higher variances. We presented these cases because we felt that it would make our test situations similar to how causal latencies behave in many real-life situations. As a result, however, the present experiment cannot address the question of whether subjects' ratings were driven more by experienced variability, more by information about expected values, or equally by both.

6. Experiment 3

In Experiments 1 and 2, deterministic causes were tested because we aimed to focus on the influence of temporal factors on singular causation judgments. However, causal strength information should also play a crucial role. A central new prediction of our model is that causal strength and temporal information are expected to interact. The goal of Experiment 3 was to explore this core property of our model.

We tested a scenario in which both causes either had a relatively high causal strength of $w_C = w_A = 0.83$ or relatively low causal strength of $w_C = w_A = 0.5$. Causal latency was manipulated by using the first pair of gamma distributions (G1 vs. G5) shown in Fig. 8. The causal latency of the target cause was either associated with G1 or G5, while the causal latency of the alternative cause always followed the complementary distribution of the pair. This combination of causal strengths and causal latencies of the potential causes leads to the prediction of an interaction effect that is shown in the middle panel of Fig. 11. As can be seen there in the left pair of bars, when the target cause has a longer causal latency than the alternative cause (i.e., its causal latency follows G5 while the causal latency of the alternative cause follows G1), the generalized model predicts that subjects should be less confident that the target cause generated the observed effect if both causes have relatively high causal strengths of 0.83 than if both causes have relatively

Table 2
Summary of the results of Experiment 2

	Condition 1	Condition 2	Condition 3	Condition 4	Condition 5
<i>Md</i>	0.90	0.80	0.50	0.30	0.10
<i>M</i>	0.91	0.69	0.54	0.37	0.11
<i>SD</i>	0.10	0.22	0.24	0.19	0.12
95% CI	[0.88, 0.94]	[0.62, 0.76]	[0.46, 0.62]	[0.31, 0.43]	[0.07, 0.15]




Fig. 10. Scatterplots depicting the correlations between predictions of the generalized model and individual singular causation judgments (left; triangles represent medians) and between model predictions and group means (right).

low causal strengths of 0.50. Conceptually, the condition in which both causes have high causal strengths of 0.83 represents a situation in which a reliably working target cause is competing with an alternative cause that not only is working very reliably but that also tends to operate quicker. According to our model and in contrast to the standard model, this combination of high causal strength and short causal latency of the competing alternative cause should make it relatively unlikely that the target cause was the singular cause of the effect. Since its high causal strength implies that the competing cause will often be strong enough to generate the effect and because its much shorter causal latency means that it almost always unfolds its causal capacity more quickly than the target cause, it should be relatively unlikely that the target cause generated the observed effect. By contrast, in the case in which both causes are less reliable and succeed in generating the effect only with a 50% chance, our model predicts that the ratings for the target cause should be higher. The reason is that the quicker competing alternative cause will relatively often fail to preempt the target cause because it is not strong enough to generate the effect, thereby leaving more room for the target cause to succeed.

Fig. 11 also shows that the predictions of our new model are reversed and align with those made by the standard model when the causal latency of the target cause is relatively short (i.e., when it follows G1; see right pair of bars). In this case, the model predicts that subjects should give higher ratings for the target cause in the condition in which both causes are relatively strong. The reason is that this condition represents a situation in which the target cause only rarely is too weak to generate the effect and in which it also tends to unfold its causal capacity more quickly than its competitor. Its causal latency advantage allows the target cause to manifest its high causal strength. By contrast, if both causes have relatively low causal strengths, ratings should be lower because the target cause will often fail to take advantage of its short causal latency.

Finally, Fig. 11 shows that the generalized model also predicts a main effect of causal latency: Causes that tend to be faster than their competitors (right pair of bars) should receive higher singular causation ratings than causes whose competitors tend to preempt

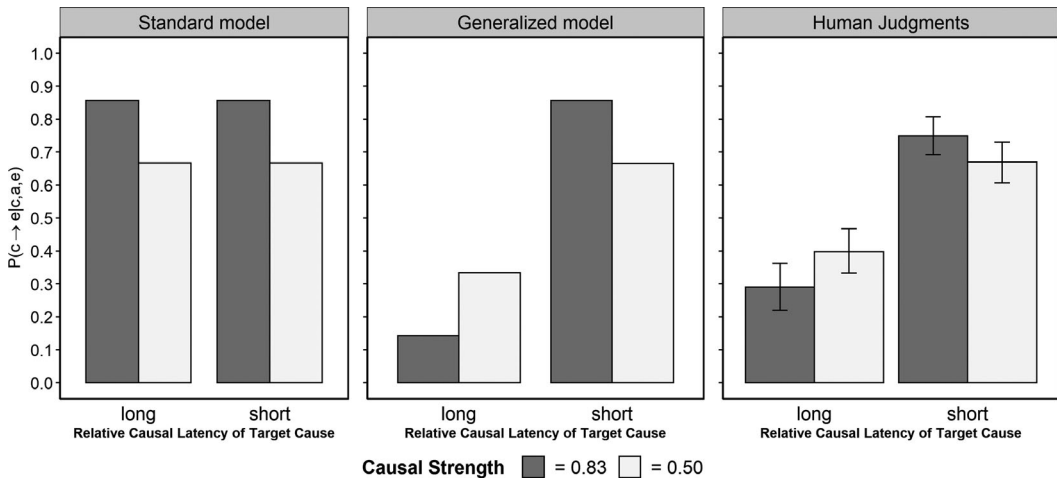


Fig. 11. Model predictions and results (bars represent mean singular causation ratings; error bars denote 95% bootstrapped CIs) of Experiment 3. The x -axes show the relative causal latency of the target cause.

them (left pair of bars). Finally, a main effect of causal strength is predicted by the standard but not by our generalized time-sensitive model for the set of parameter values we used. It is important to note, though, that the absence of a predicted main effect of causal strength does not mean that causal strength information is per se considered to be less relevant than latency information by the generalized model. The absence of a predicted main effect of causal strength is a result of the particular parameter value combination we chose in this study. Specifically, it is mainly due to the fact that the target and alternative cause have identical causal strength, whereas they differ with respect to causal latency. We used this combination of parameter values for which our generalized model predicts an interaction but no main effect of causal strength because it gives us a strong test of our new model.

6.1. Methods

6.1.1. Participants

One hundred and sixty subjects ($M_{\text{age}} = 38.14$, $SD_{\text{age}} = 12.20$, 116 female, 44 male) were recruited via Prolific (www.prolific.ac). This sample size ($n = 40$ subjects per condition) allows us to detect medium effects of $d = 0.60$ for simple group comparisons on the basis of directed t tests with more than 80% test power. The inclusion criteria were the same as in the previous studies. Subjects were paid £1.20 for their successful participation.

6.1.2. Design, materials, and procedure

A demo video of this experiment is available under <https://osf.io/hdwx3/>. Subjects were randomly assigned to one of four conditions ($n = 40$) that resulted from a 2 (causal strength of the two causes: $w_C = w_A = 0.83$ vs. $w_C = w_A = 0.50$) \times 2 (relative causal

latency of the target cause: long vs. short) between-subjects design. The following factors were counterbalanced between subjects: the order of presentation of the two causes in the learning phase (i.e., tower “North” first vs. tower “South” first); which cause had the shorter/longer relative causal latency; the target cause to which the main test question referred; the orientation of the rating scale belonging to the main test question (high confidence on the left vs. on the right side of the scale).

The materials and the procedure were largely identical to those of Experiments 1 and 2 with the following exceptions: First, information about the probabilistic causal nature of the towers was added to the instructions by informing subjects that pigeons might get lost on their way to the palace and that it would therefore be important to study the pigeons’ arrival rates. Secondly, the learning task was modified so that causal strength and causal latency information were conveyed together.

Unlike in Experiment 2, subjects observed 24 pigeons per tower in this experiment. Our goal was to ensure that participants observed sufficiently many successful pigeons to be able to learn the causal latencies. Depending on the causal strength condition, subjects observed 20 or 12 successful pigeons per tower. The flight durations corresponded to 12 versus 20 percentiles of the respective causal latency distributions. Whenever a pigeon failed to reach the palace, the words “pigeon probably lost” were displayed 5 s after the pigeon had been sent out. This delay corresponded to the 99.9th percentile of the slower latency distribution. The singular causation test question was the same as in Experiments 1 and 2.

Additionally, on a separate screen subjects were asked to estimate the causal strengths of the two causes because we wanted to control for the possibility that subjects’ causal strength representations of the two causes might have been influenced by the causes’ causal latencies. The new model considers causal strength and causal latency to be two independent properties of causes. Yet subjects might not have kept causal strength and latency information apart. They might have, for instance, represented a relatively slow cause as less powerful than a relatively fast cause. To assess subjects’ causal strength representations, they were asked the following question for each tower, with the order of presentation being randomized: “Based on what you have learned: how many out of 10 letter pigeons sent from Tower ‘North’ (‘South’) would make it to the palace?” Ratings were provided on an 11-point scale (0–10).

6.2. Results and discussion

The results are summarized in the right panel of Fig. 11. Singular causation ratings followed the predictions of the new model, whereas the standard model does not explain the results. As for the predicted interaction, Fig. 11 shows that when the target cause had a relatively long causal latency, ratings in the low causal strength condition were higher ($M = 0.40$, $SD = 0.23$, $Md = 0.40$, 95% CI [0.33, 0.47]) than those in the high causal strength condition ($M = 0.29$, $SD = 0.26$, $Md = 0.20$, 95% CI [0.21, 0.37]). The reversed pattern was obtained when the target cause had relatively short causal latency ($M = 0.67$, $SD = 0.20$, $Md = 0.70$, 95% CI [0.61, 0.73] when causal strength was 0.50 vs. $M = 0.75$, $SD = 0.19$, $Md = 0.80$, 95% CI [0.69, 0.81] when causal strength was 0.83). A planned

contrast testing the predicted interaction was significant, $t(156) = 2.68$, $p < .01$, $d = 0.32$. However, Fig. 11 also shows that the simple effects qualifying the interaction were smaller than predicted by the generalized model ($t(156) = 2.17$, $p = .016$ one-sided, $d = 0.44$ for the condition in which the target cause had a relatively long causal latency vs. $t(156) = 1.62$, $p = .05$ one-sided, $d = 0.41$ for the condition in which the target cause's causal latency was relatively short).

Figure 11 also shows a main effect of the relative causal latency of the target cause. A planned contrast testing this predicted main effect was significant, $t(156) = 10.47$, $p < .001$, $d = 1.62$. Finally, the main effect of causal strength, which is predicted by the standard but not by the generalized model, was not found.

The causal strength ratings subjects made are summarized in Fig. 12. Causal strength estimates were not distorted by the different causal latencies of the two causes ($M = 5.16$, $SD = 1.07$, $Md = 5.00$, 95% CI [4.92, 5.40] vs. $M = 5.31$, $SD = 1.18$, $Md = 5.00$, 95% CI [5.05, 5.57], for the condition with causal strength = 0.50 and $M = 7.58$, $SD = 1.75$, $Md = 8.00$, 95% CI [7.19, 7.96] vs. $M = 7.66$, $SD = 1.86$, $Md = 8.00$, 95% CI [7.25, 8.08] for the condition with causal strength = 0.83). The causal strength ratings were close to the normative values of 0.83 and 0.50. A mixed ANOVA with relative causal latency of the target cause as a within-subject factor and causal strengths of the causes as between-subjects factor only yielded a significant main effect for causal strength, $F(1, 158) = 122.68$, $p < .001$, $d = 1.6$. This main effect confirms that causal strength ratings were solely determined by causal strength and not by causal latency. These results rule out that subjects' singular causation ratings were influenced by differences in the perceived causal strengths of the causes.

7. Experiment 4

Although the pattern of singular causation judgments in Experiment 3 was captured well by the new model, the predicted interaction effect turned out to be relatively weak. Since the predicted interaction follows from a core principle of the new model, we wanted to replicate the results in a fourth, pre-registered experiment with a selected set of causal power parameters that should lead to a larger interaction effect. The same causal latency distributions were used as in Experiment 3, but this time causes were contrasted that either had a causal strength of 0.93 or 0.40. The predictions are shown in Fig. 13. The experiment was pre-registered at the Open Science Framework (OSF; <https://osf.io/>). The pre-registration can be accessed under <https://doi.org/10.17605/OSF.IO/N93BU>.

At the end of Experiment 3, subjects had been asked to estimate the causal strength of the two causes because it was important to control for the possibility that subjects' causal strength representations are influenced by the observed differences in causal latencies. We found that this was not the case. However, it is possible that there is an influence in the other direction. Differences in causal strengths might have an impact on causal latency representations. This is not unlikely in the paradigm we used because trials in which a cause failed to generate the effect ended after a duration that corresponded to the 99.9th percentile of the

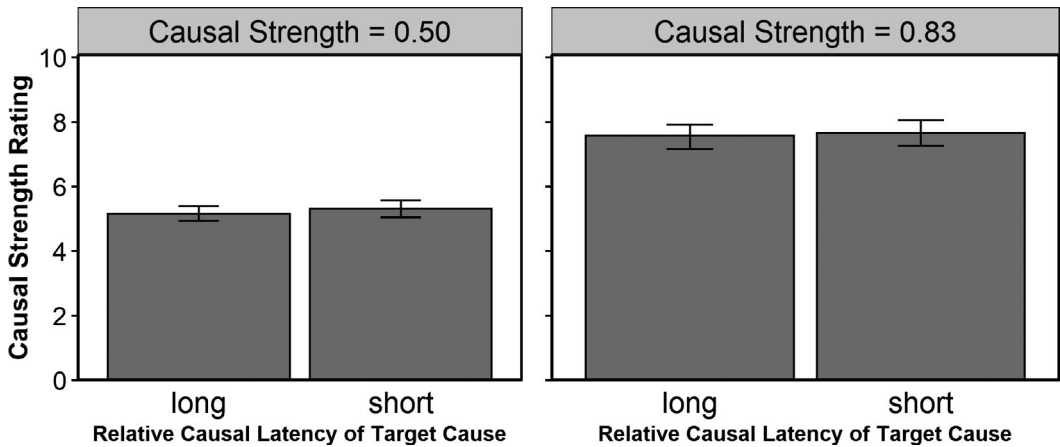


Fig. 12. Results (bars represent means; error bars denote 95% bootstrapped CIs) for the causal power question in Experiment 3. The left panel shows the ratings for the condition in which both causes had a causal power of 0.50. The right panel shows the ratings for the condition in which both causes had a causal power of 0.83.

slower latency distribution (G5). Consequently, subjects in the low causal strength condition spent more time on the task than subjects in the high causal strength condition. In the present study, we therefore included a control question assessing how much subjects' causal latency representations were affected by causal strength information.

7.1. Methods

7.1.1. Participants

Three hundred and eighty-four subjects ($M_{\text{age}} = 34.25$, $SD_{\text{age}} = 11.38$, 182 female, 200 male, two subjects did not want to report their gender) were recruited via Prolific (www.prolific.ac). Subjects were at least 18 years old, had at least an A-level degree, were native English speakers, and had an approval rate concerning participation in previous of at least 90%. They were paid £1.20 for participation. The sample size calculation was done with the program G*Power 3.1 (Faul, Erdfelder, Lang, & Buchner, 2007) and based on the effect sizes found in a set of pilot studies conducted after the third experiment and prior to the pre-registration of the new study. Both Experiment 3 and these pilot studies indicated that the smallest difference between the means is to be expected in the condition in which the target cause has a relatively short causal latency (see right pair of bars in Fig. 11). We focused on this smallest difference when we planned the present study. We calculated the sample size such that independent-samples t tests (one-sided) testing the two mean differences involved in the predicted interaction effect would detect medium effects of $d = 0.5$ ($\alpha = 0.05$) with about 95% probability. This conservative criterion also implies that an ANOVA would detect the predicted interaction effect with even higher probability.

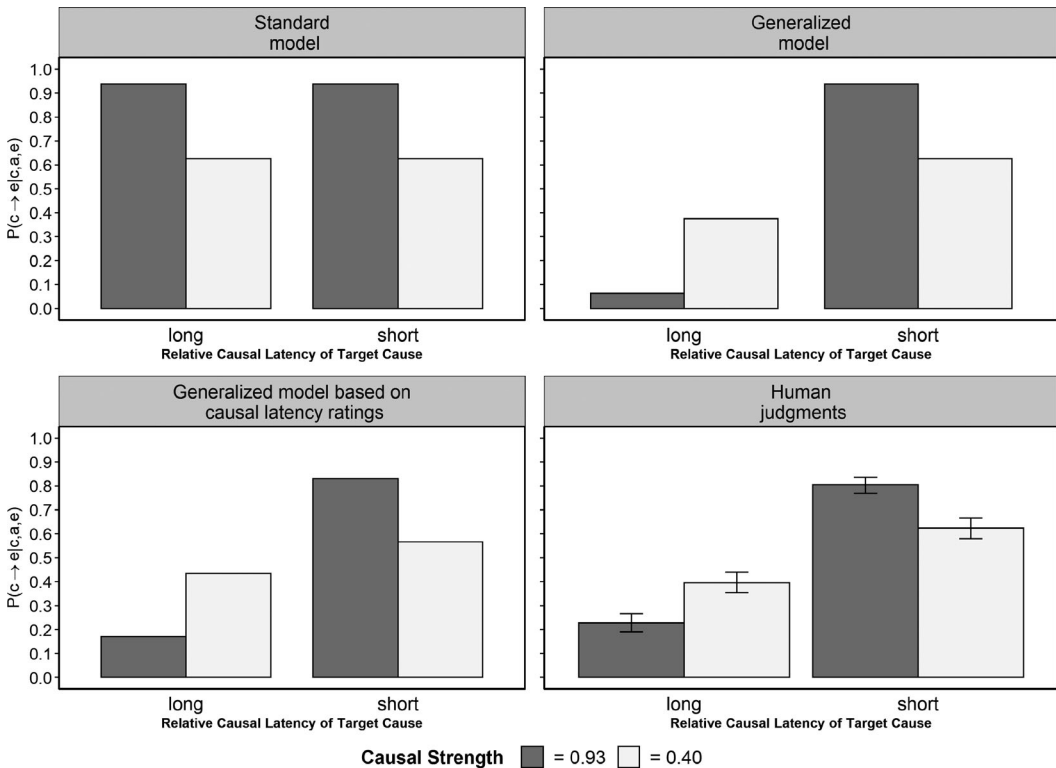


Fig. 13. Model predictions and results (bars represent mean singular causation ratings; error bars denote 95% bootstrapped CIs) of Experiment 4.

7.1.2. Design, materials, and procedure

A demo video of the experiment is available under <https://osf.io/7x6tu/>. The study design was identical to the one of Experiment 3. Subjects were randomly assigned to one of four conditions ($n = 96$) resulting from a 2 (causal strengths of the two causes: $w_C = w_A = 0.93$ vs. $w_C = w_A = 0.40$) \times 2 (relative causal latency of target cause: long vs. short) between-subjects design.

Since this experiment contrasted causal strength values of 0.93 and 0.40, the learning phase differed from the one in Experiment 3. In the condition in which the causes had causal strengths of 0.93, subjects learned that each cause was successful in generating the effect in 28 out of 30 attempts. In the low causal strength condition, each cause generated the effect on 12 out of 30 occasions.

After subjects had answered the singular causation query, they were asked a control question referring to the causal latencies of the two causes. Subjects were asked to imagine an occasion on which both towers simultaneously sent a pigeon. They then should indicate on an 11-point scale with endpoints “definitely the pigeon from Tower ‘North’” and “definitely the pigeon from Tower ‘South’” (the midpoint was labeled “50:50”)

which pigeon they think would arrive at the palace first. Answers to this question allowed us to estimate participants' α parameter. The orientation of the scale (whether tower North was on the left or right side of the scale) was randomly varied between subjects.

7.2. Results and discussion

The results are summarized in the fourth panel of Fig. 13, which shows that the replication was successful. The results of a 2 (causal power of the two causes: 0.93 vs. 0.40) \times 2 (relative causal latency of target cause: long vs. short) factorial ANOVA confirmed that all predicted effects were significant. There was a main effect of "relative causal latency of target cause", $F(1, 380) = 381.02$, $p < .001$, $d = 2.0$, confirming that singular causation ratings were overall higher when the target cause had a shorter causal latency than its competitor. Importantly, the interaction between causal strength and relative causal latency of target cause was also significant, $F(1, 380) = 71.19$, $p < .001$, $d = 0.87$. Contrast analyses testing the simple effects qualifying the interaction showed that the difference between high and low causal strength was significant in both causal latency conditions. In the condition in which the target cause's causal latency was longer than the one of the alternative cause, subjects gave higher singular causation ratings when causal strength was 0.40 ($M = 0.39$, $SD = 0.22$, $Md = 0.40$, 95% CI [0.35, 0.44]) than when causal strength was 0.93 ($M = 0.23$, $SD = 0.20$, $Md = 0.20$, 95% CI [0.19, 0.27]), $t(380) = 5.74$, $p < .001$, $d = 0.85$. In the condition in which the target cause's causal latency was shorter than the one of the alternative cause, the reverse pattern was found ($M = 0.62$, $SD = 0.22$, $Md = 0.70$, 95% CI [0.58, 0.67] for causal strength = 0.40; $M = 0.81$, $SD = 0.17$, $Md = 0.80$, 95% CI [0.77, 0.84] for causal strength = 0.93), $t(380) = 6.20$, $p < .001$, $d = 1.05$. As predicted by our model, there was no main effect of causal strength, $F(1, 380) < 1$. Furthermore, the effects were overall larger than the ones we observed in Experiment 3, which indicates that subjects were indeed sensitive to the causal strength difference. In Experiment 3, the size of the interaction effect was $d = 0.32$, whereas it was $d = 0.87$ in the present study.

We next consider the results for the causal latency control question subjects answered toward the end of the study. We transformed subjects' ratings across the different counterbalancing factors so that high ratings are associated with the cause with the shorter causal latency. We found that ratings were overall high in both causal strength conditions, confirming that subjects tended to correctly identify the cause with the shorter causal latency. However, we also found that our causal strength manipulation seems to have led to a small distortion: In the condition in which both causes were strong, subjects were more likely to make ratings associated with the cause with the shorter causal latency ($M = 0.877$, $SD = 0.169$, $Md = 0.90$, 95% CI [0.852, 0.90]) than in the condition in which the causes were weak ($M = 0.764$, $SD = 0.254$, $Md = 0.90$, 95% CI [0.728, 0.80]). One explanation for this difference might be that trials during the learning phase in which a cause failed to generate the effect in fact lasted longer than successful trials (see procedure). Subjects in the low causal strength condition saw more of

these trials, and it is possible that these trials contaminated their causal latency representations. (We will consider another possible explanation for this effect in Section 8.)

The model predictions shown in Panel 2 of Fig. 13 were obtained with α values that were computed based on the objective causal latency parameters. However, we can also use the responses subjects gave for the latency control question to estimate α . We therefore also examined how using these subjective values for the α parameter (0.877 and $1 - 0.877$ for high vs. low alpha in the high-causal power case; 0.764 and $1 - 0.764$ for high vs. low alpha in the low-power case) would change the model predictions. The model predictions relying on the subjective α values are depicted in the third panel in Fig. 13. It can be seen that, because the subjective α values were less extreme than the objective ones, the model predictions for all the conditions became slightly less extreme and moved a bit closer to the midpoint of the scale. Comparing Panels 3 and 4 in Fig. 13, it can be seen that using participants' subjective α values led to an even closer fit between model predictions and subjects' singular causation judgments than when using the objective parameter values (Panel 2). The results of this experiment demonstrate again the validity of the new model.

8. General discussion

The present article argues that assessments of singular causation require the integration of two different types of information: First, causal strength of the potential causes is important. If an outcome can be explained by either of two potential causes, then the present stronger cause is more likely to be the singular cause than the weaker present cause. The standard causal power theory of attribution by Cheng and Novick (2005) provides a formal analysis of how the causal strengths of the potential causes should be weighed against each other to reach a decision in a singular causation scenario. Relying on the generalized power model of causal attribution by Stephan and Waldmann (2018), we here argued and demonstrated that considering only causal strength is not enough, however. The second crucial type of information is information about the potential causes' temporal relation. Without temporal information, it would be impossible to distinguish situations in which a sufficiently strong cause actually succeeds in bringing about the effect from those in which a likewise sufficient competitor manages to preempt the target cause. We therefore proposed a formal model that incorporates the probability of preemption to estimate how likely it is that the target cause actually caused the target outcome. We showed that the probability of preemption can be estimated by the combination of causal strength and temporal information about cause onset times and causal latencies. The results of the reported experiments demonstrate that the new model makes accurate predictions of subjects' singular causation judgments. Interestingly, the model was accurate even in the computationally quite demanding situations in which its parameters were probabilistic (Experiments 3 and 4).

8.1. Limitations and avenues for future studies

8.1.1. Modeling more complex types of situations

The type of test situations we investigated represents only a subset of possible situations. For example, we here focused on rather simple situations in which there were only two potential causes of an observed effect. Subjects' singular causation judgments were well predicted by our model in these cases. However, often there are more than two potential causes of an effect. Our model can also be applied to these cases, and it would be interesting to see how subjects' performance changes in these more complicated conditions in which tracking and integrating causal strength and temporal information about multiple causes represent a cognitively quite challenging task.

Another noteworthy characteristic of our test scenarios was that all potential causes of the effect were actually observed. In many real-world situations, however, reasoners observe only a subset of the potential causes and will therefore be confronted with uncertainty concerning the presence of further alternative causes. Although we have not considered situations with unobserved background causes here, the new model can principally be applied to such situations, too. As Cheng (1997) has shown, when alternative causes of the effect are unobserved and independent of the target cause, the causal strength parameter w_C of the target cause can be estimated by $\Delta P / (1 - P[e|\neg c])$ and the base rates and strengths of the potential alternative causes is estimated by $P(e|\neg c)$. The causal strength parameters of the model can therefore be estimated in these cases. A novel question is how the α parameter can be estimated in such situations. One idea is that the computation of α can be realized by carrying out multiple steps. First, information about the temporal distribution of the effect in the absence of the target cause needs to be considered. As Bramley et al. (2018) have shown in their article on general causal structure learning, the temporal distribution of the effect in the target cause's absence can be modeled with exponential distributions.

Next, it needs to be analyzed how the distribution changes in the presence of the target cause. The observed causal delay distribution in the presence of the target cause would be a mixture of the gamma distribution of the target cause's causal latency and the previously measured exponential background distribution. In a next step, the previously gained knowledge about the background distribution can be used to decompose the mixed distribution into its elements. Once the causal latency distribution of the target cause has been extracted from the mixed distribution, the same algorithm that we used to estimate α in the present studies can be applied. We plan to test such situations in future experiments. Again, we expect that target causes with a short causal latency should more likely be viewed as singular causes than otherwise identical causes with longer causal latency, as the latter are more likely to be intercepted by unobserved background causes (see Lagnado & Speekenbrink, 2010).

Furthermore, the present studies tested only static target situations. Subjects were asked to make singular causation judgments in a specific scenario based on knowledge about strength, delays, and onsets that was acquired in separate learning situations. The

structure of the target situation thus resembled a classical crime scenario in which the detective enters the scenery and tries to identify the perpetrator after all events had unfolded. We focused on this type of test situation because it is the type of situation that was studied in previous related studies (e.g., Cheng & Novick, 2005; Holyoak, Lee, & Lu, 2010; Meder, Mayrhofer, & Waldmann, 2014). In many real-life situations, however, singular causation queries arise in dynamic contexts. A crucial feature of dynamic target situations is that the onsets and delays are experienced during the observation phase prior to the singular causation query. This way, the interaction between prior beliefs about causal strength and temporal parameters could be compared to the actually observed sequence of events. Imagine a target cause with an average causal strength of $w_C = 0.90$ and an expected causal delay of 20 min (e.g., a dose of ibuprofen). Now imagine a reasoner observes that the relief from pain occurs after 2 hours. The unexpected long delay may be taken as diagnostic evidence that the generally positive causal strength of the drug was zero on this occasion, and therefore might weaken the belief that the drug was the singular cause of the effect. How reasoners make singular causation judgments in dynamic target situations and how their judgments can be modeled in this case is an interesting question for future investigations (for a first pass on how this online attribution problem could be handled, see Bramley, Mayrhofer, Gerstenberg, & Lagnado, 2017).

8.1.2. The influence of other psychological factors

Since we here proposed a rational computational model of singular causation judgments, we have largely left unaddressed other plausible psychological factors that might constrain subjects' performance. Two of these factors that may have influenced subjects' answers in our tasks are working memory constraints and psychophysical effects. There are, however, also less obvious factors that might explain deviations of reasoners' performance from the model predictions. For example, one interesting effect found in previous studies is the so-called temporal binding effect.⁵ The effect describes the phenomenon that knowledge about the existence of a general causal relationship between two factors can diminish the perceived temporal distance between cause and effect (Buehner, 2012; Faro, Leclerc, & Hastie, 2005; Faro, McGill, & Hastie, 2010; Marsh & Ahn, 2009; see also Bechlivanidis & Lagnado, 2016, for the even more intriguing phenomenon that causal knowledge might sometimes even change the perception of the order in which events occurred). One result in our studies that is in line with the temporal binding effect are the causal latency ratings we assessed in Experiment 4. We found that subjects tended to report a shorter causal latency for the cause that had higher causal strength, even though the effect was small. To which extent this finding is actually due to a temporal binding effect or due to other features of our experimental task is an interesting question open for future research.

8.1.3. The role of causal mechanism information

Another question that we left unaddressed in the present paper concerns the role of causal mechanism information in the assessment of singular causation. Causal mechanism information has been identified as an important type of information in causal reasoning

(see Cartwright, 2017; Danks, 2017; Johnson & Ahn, 2017, for overviews). Different studies have shown that reasoners regard mechanism information as particularly relevant when making causal attributions (Ahn, Kalish, Medin, & Gelman, 1995; Johnson & Keil, 2018).

An interesting step for future work would be to experimentally study the role of mechanism information in singular causation judgments and to incorporate this factor in our computational model. In a first approximation, causal mechanisms can be modeled in causal Bayes nets as causal chains or causal networks in which the target cause and effect are connected indirectly by one or more intermediate variables. It is possible to extend our model to incorporate mechanism information by adding variables mediating between the target cause and effect. However, an additional complexity arises because mechanisms are not only chains or networks of interconnected variables but also extend in time. Causal mechanism information not only constrains the causal strength parameters in the model but also the temporal parameters attached to the different causal relations. For example, a causal mechanism connecting a cause and an effect might involve multiple causal paths which might differ not only with respect to their causal strengths but also with respect to their causal latencies. Take the example of a coroner who wants to find out whether a victim displaying an abdominal gunshot wound actually died because of the gunshot. The causal strength with which bullets kill their victims is surely quite high on average, but also fairly variable. Similarly, bullets typically bring death to their victims relatively quickly, but not always. Depending on the organs that are damaged, both the probability of dying from a gunshot and the latency can be high or low. If the bullet hit the victim's heart, one will be more confident in a causal link between gunshot and death than in a case in which the bullet left all organs intact. Bullets that stop their victims' hearts are not only more effective, they also bring death so quickly that there is less room for alternative causes to strike (e.g., an assassin might first have administered a lethal dose of poison but then decided to shoot the victim). Testing the influence of causal mechanism information on singular causation judgments will be an interesting avenue for future studies.

8.2. Conclusion

In the present research, we have shown that singular causal judgments are influenced by both causal strength and temporal information. We have proposed a computational model that integrates these factors and thereby explicitly considers the role of causal pre-emption. Yet our research only represents a first step. Future versions of our model need to consider more complex situations in which events dynamically unfold in time and are mediated by more or less complex networks of mechanisms. It would also be interesting to widen the range of test questions our computational model can address. The focus here was on singular causation queries in which the potential causes and the target effect were all observed. It would be interesting to also study situations in which some or all of the potential causes remain unobserved.

Acknowledgments

This research was supported by Deutsche Forschungsgemeinschaft (DFG) Grants WA 621/24-1 and MA 6545/1-2, the latter as part of the priority program “New Frameworks of Rationality” (SPP 1516). A preliminary report of Experiments 2 and 3 was presented at the 2018 Annual Conference of the Cognitive Science Society in Madison, Wisconsin. We thank Patricia Cheng, York Hagmayer, Keith Holyoak, Hongjing Lu, and Alex Wiegmann for helpful discussions about this project. We thank Louisa Reins for her assistance in the designing and implementation of the experiments.

Declaration of Interest

None.

Open Research Badges



This article has earned Open Data and Open Materials Badges. Data and materials are available at https://osf.io/5yvs4/?view_only=c574902bee2740e9b49cd872c7ca2f89

Notes

1. The psychological validity of this independence assumption has been confirmed in several empirical studies that investigated default assumptions of human learners (Buehner & Cheng, 1997; Cheng, 1997; Cheng & Lu, 2017; Griffiths & Tenenbaum, 2005; Liljeholm & Cheng, 2007; Lu, Yuille, Liljeholm, Cheng, & Holyoak, 2008; Meder & Mayrhofer, 2017; Meder, Mayrhofer, & Waldmann, 2014).
2. The same sampling algorithm can be used to estimate α for fixed causal latencies. However, it is easy to see that α will either be 1 or 0 in such situations depending on whether the sum of $t_{(a \rightarrow e)}$ and Δ_t is smaller than $t_{(c \rightarrow e)}$ or not.
3. Empirical studies that explicitly investigated this type of situation have yielded mixed results (see Sloman & Lagnado, 2015, for an overview). Some studies (e.g., Spellman & Kincannon, 2001) found that people seem to regard symmetrically overdetermining causes as highly causal. Others found that such cases are viewed as less causal than cases of preemption (e.g., Gerstenberg, Goodman, Lagnado, & Tenenbaum, 2015).
4. One may wonder whether subjects' responses were purely causal or, since subjects read that the people of Extonia wanted to decorate the crew who caused the alarm, whether they also involved considerations about praise and blame (see Samland &

Waldmann, 2016). While this is a possibility, we do not think that it undermines our results because (a) praise and blame are closely linked to causality, and (b) because the phrasing of the test situation was kept constant in the present and all the following studies.

5. We thank one of the reviewers for pointing out this possibility.

References

- Ahn, W.-K., Kalish, C. W., Medin, D. L., & Gelman, S. A. (1995). The role of covariation versus mechanism information in causal attribution. *Cognition*, *54*, 299–352.
- Bechlivanidis, C., & Lagnado, D. A. (2016). Time reordered: Causal perception guides the interpretation of temporal order. *Cognition*, *146*, 58–66.
- Bramley, N., Gerstenberg, T., Mayrhofer, R., & Lagnado, D. A. (2018). The role of time in causal structure learning. *Journal of Experimental Psychology: Learning, Memory & Cognition*, *44*, 1880–1910.
- Bramley, N., Mayrhofer, R., Gerstenberg, T., & Lagnado, D. A. (2017). Causal learning from interventions and dynamics in continuous time. In G. Gunzelmann, A. Howes, T. Tenbrink, & E. Davelaar (Eds.), *Proceedings of the 39th Annual Conference of the Cognitive Science Society* (pp. 150–155). Austin, TX: Cognitive Science Society.
- Buehner, M. J. (2012). Understanding the past, predicting the future: Causation, not intentional action, is the root of temporal binding. *Psychological Science*, *23*, 1490–1497.
- Buehner, M. J. (2017). Space, time, and causality. In M. R. Waldmann (Ed.), *The Oxford handbook of causal reasoning* (pp. 549–564). New York: Oxford University Press.
- Buehner, M. J., & Cheng, P. W. (1997). Causal induction: The Power PC theory versus the Rescorla-Wagner model. In M. Shafto & P. Langley (Eds.), *Proceedings of the 19th Annual Conference of the Cognitive Science Society*, (pp. 55–60). Hillsdale, NJ: Lawrence Erlbaum Associate.
- Buehner, M. J., & McGregor, S. (2006). Temporal delays can facilitate causal attribution: Towards a general timeframe bias in causal induction. *Thinking & Reasoning*, *12*, 353–378.
- Cartwright, N. (2017). Single case causes: What is evidence and why. In H. K. Chao & J. Reiss (Eds.), *Philosophy of science in practice* (pp. 11–24). Cham, Switzerland: Springer.
- Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, *104*, 367–405.
- Cheng, P. W., & Lu, H. (2017). Causal invariance as an essential constraint for creating a causal representation of the world: Generalizing the invariance of causal power. In M. R. Waldmann (Ed.), *The Oxford handbook of causal reasoning* (pp. 65–84). New York: Oxford University Press.
- Cheng, P. W., & Novick, L. R. (1991). Causes versus enabling conditions. *Cognition*, *40*, 83–120.
- Cheng, P. W., & Novick, L. R. (2005). Constraints and nonconstraints in causal learning: Reply to White (2005) and to Luhmann and Ahn (2005). *Psychological Review*, *112*, 694–706.
- Danks, D. (2017). Singular causation. In M. R. Waldmann (Ed.), *The Oxford handbook of causal reasoning* (pp. 201–215). New York: Oxford University Press.
- Faro, D., Leclerc, F., & Hastie, R. (2005). Perceived causality as a cue to temporal distance. *Psychological Science*, *16*, 673–677.
- Faro, D., McGill, A. L., & Hastie, R. (2010). Naïve theories of causal force and compression of elapsed time judgments. *Journal of Personality and Social Psychology*, *98*, 683–701.
- Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2007). G* power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, *39*, 175–191.
- Gerstenberg, T., Goodman, N. D., Lagnado, D. A., Tenenbaum, J. B. (2015). How, whether, why: Causal judgments as counterfactual contrasts. In D. C. Noelle, R. Dale, A. S. Warlaumont, J. Yoshimi, T.

- Matlock, C. D., Jennings, & P. P. Maglio (Eds.), *Proceedings of the 37th Annual Conference of the Cognitive Science Society* (pp. 782–787). Austin, TX: Cognitive Science Society.
- Griffiths, T. L., & Tenenbaum, J. B. (2005). Structure and strength in causal induction. *Cognitive Psychology*, *51*, 334–384.
- Hagmayer, Y., & Waldmann, M. R. (2002). How temporal assumptions influence causal judgments. *Memory & Cognition*, *30*, 1128–1137.
- Halpern, J. Y. (2016). *Actual causality*. Cambridge, MA: MIT Press.
- Halpern, J. Y., & Hitchcock, C. (2015). Graded causation and defaults. *The British Journal for the Philosophy of Science*, *66*, 413–457.
- Hart, H. L. A., & Honoré, T. (1959/1985). *Causation in the law*. Oxford, UK: Oxford University Press.
- Hitchcock, C. (2001). The intransitivity of causation revealed in equations and graphs. *The Journal of Philosophy*, *98*, 273–299.
- Hitchcock, C. (2007). Prevention, preemption, and the principle of sufficient reason. *The Philosophical Review*, *116*, 495–532.
- Hitchcock, C. (2009). Causal modelling. In H. Beebe, C. Hitchcock, & P. Menzies (Eds.), *The Oxford handbook of causation* (pp. 299–314). New York: Oxford University Press.
- Hitchcock, C., & Knobe, J. (2009). Cause and norm. *The Journal of Philosophy*, *106*, 587–612.
- Holyoak, K. J., Lee, H. S., & Lu, H. (2010). Analogical and category-based inference: A theoretical integration with Bayesian causal models. *Journal of Experimental Psychology: General*, *139*, 702–727.
- Icard, T. F., Kominsky, J. F., & Knobe, J. (2017). Normality and actual causal strength. *Cognition*, *161*, 80–93.
- Johnson, S. G., & Ahn, W.-K. (2017). Causal mechanisms. In M. R. Waldmann (Ed.), *The Oxford handbook of causal reasoning* (pp. 127–146). New York: Oxford University Press.
- Johnson, S. G., & Keil, F. C. (2018). Statistical and mechanistic information in evaluating causal claims. In T. T. Rogers, M. Rau, X. Zhu, & C. W. Kalish (Eds.), *Proceedings of the 40th Annual Conference of the Cognitive Science Society* (pp. 618–623). Austin, TX: Cognitive Science Society.
- Kominsky, J. F., Phillips, J., Gerstenberg, T., Lagnado, D. A., & Knobe, J. (2015). Causal superseding. *Cognition*, *137*, 196–209.
- Lagnado, D. A., & Gerstenberg, T. (2017). Causation in legal and moral reasoning. In M. R. Waldmann (Ed.), *The Oxford handbook of causal reasoning* (pp. 565–602). New York, NY: Oxford University Press.
- Lagnado, D. A., Gerstenberg, T., & Zultan, R. (2013). Causal responsibility and counterfactuals. *Cognitive Science*, *37*, 1036–1073.
- Lagnado, D. A., & Speekenbrink, M. (2010). The influence of delays in real-time causal learning. *The Open Psychology Journal*, *3*, 184–195.
- Lewis, D. (1973). Causation. *The Journal of Philosophy*, *17*, 556–567.
- Liljeholm, M., & Cheng, P. W. (2007). When is a cause the “same”? Coherent generalization across contexts. *Psychological Science*, *18*, 1014–1021.
- Lombrozo, T., & Vasilyeva, N. (2017). Causal explanation. In M. R. Waldmann (Ed.), *The Oxford handbook of causal reasoning* (pp. 415–432). New York: Oxford University Press.
- Lu, H., Yuille, A. L., Liljeholm, M., Cheng, P. W., & Holyoak, K. J. (2008). Bayesian generic priors for causal learning. *Psychological Review*, *115*, 955–982.
- 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*, 334–352.
- Meder, B., & Mayrhofer, R. (2017). Diagnostic causal reasoning with verbal information. *Cognitive Psychology*, *96*, 54–84.
- Meder, B., Mayrhofer, R., & Waldmann, M. R. (2014). Structure induction in diagnostic causal reasoning. *Psychological Review*, *121*, 277–301.
- Menzies, P. (1996). Probabilistic causation and the pre-emption problem. *Mind*, *105*, 85–117.
- Novick, L. R., & Cheng, P. W. (2004). Assessing interactive causal influence. *Psychological Review*, *111*, 455–485.

- Paul, L. A., & Hall, E. J. (2013). *Causation: A user's guide*. Oxford, UK: Oxford University Press.
- Pearl, J. (1988). *Probabilistic reasoning in intelligent systems: Networks of plausible inference*. San Francisco, CA: Morgan Kaufmann.
- Pearl, J. (2000). *Causality: Models, reasoning and inference*. Cambridge, UK: Cambridge University Press.
- Phillips, J., Luguri, J. B., & Knobe, J. (2015). Unifying morality's influence on non-moral judgments: The relevance of alternative possibilities. *Cognition*, *145*, 30–42.
- Rottman, B. M. (2017). The acquisition and use of causal structure knowledge. In M. R. Waldmann (Ed.), *The Oxford handbook of causal reasoning* (pp. 85–114). New York: Oxford University Press.
- Rottman, B. M., & Hastie, R. (2014). Reasoning about causal relationships: Inferences on causal networks. *Psychological Bulletin*, *140*, 109–139.
- Samland, J., & Waldmann, M. R. (2016). How prescriptive norms influence causal inferences. *Cognition*, *156*, 164–176.
- Schaffer, J. (2003). Overdetermining causes. *Philosophical Studies*, *114*, 23–45.
- Shanks, D. R., Pearson, S. M., & Dickinson, A. (1989). Temporal contiguity and the judgement of causality by human subjects. *The Quarterly Journal of Experimental Psychology*, *41*, 139–159.
- Shortle, J. F., Thompson, J. M., Gross, D., & Harris, C. M. (2018). *Fundamentals of queueing theory*. Hoboken, NJ: John Wiley & Sons.
- Sloman, S. (2005). *Causal models: How people think about the world and its alternatives*. New York: Oxford University Press.
- Sloman, S., & Lagnado, D. A. (2015). Causality in thought. *Annual Review of Psychology*, *66*, 223–247.
- Spellman, B. A., & Kincannon, A. (2001). The relation between counterfactual (but for) and causal reasoning: Experimental findings and implications for jurors' decisions. *Law and Contemporary Problems: Causation in Law and Science*, *64*, 241–264.
- Spirtes, P., Glymour, C., & Scheines, R. (1993). *Causation, prediction, and search*. New York: Springer-Verlag.
- Stephan, S., & Waldmann, M. R. (2018). Preemption in singular causation judgments: A computational model. *Topics in Cognitive Science*, *10*, 242–257.
- Strevens, M. (2008). *Depth: An account of scientific explanation*. Cambridge, MA: Harvard University Press.
- Waldmann, M. R. (Ed.). (2017). *The Oxford handbook of causal reasoning*. New York: Oxford University Press.