

Assessing Singular Causation: The Role of Causal Latencies

Simon Stephan (sstepha1@gwdg.de) Ralf Mayrhofer (rmayrho@gwdg.de)

Michael R. Waldmann (michael.waldmann@bio.uni-goettingen.de)

Department of Psychology, University of Göttingen,
Gosslerstr. 14, 37073 Göttingen, Germany

Abstract

Singular causation queries require an assessment of whether a singular co-occurrence of two events c and e was causal or simply coincidental. The current study builds on our previous research (Stephan & Waldmann, 2018) in which we proposed a computational model of singular causation judgments. The model highlights that singular causation judgments need to take into account the power of the target cause C and of alternative causes A , as well as the possibility of preemption. What was missing was a detailed model allowing us to estimate the probability of preemption of a target cause by the alternative causes. The present research fills this gap by elaborating the temporal assumptions that might enter assessments of singular causation. We focus on assumptions about temporal precedence between target and alternative causes, with a specific focus on assumptions about causal latency. We report the results of two new experiments supporting the model.

Keywords: singular causation; causal attribution; preemption; time; causal reasoning; computational modeling

Causal representations support various inferences. They enable us to make predictions, to form diagnoses, or to make judgments about singular causation (Waldmann, 2017). In the present research our goal is to develop and test a model explaining judgments about singular causation. How does a person come to believe, for example, that it was the medicine she took that caused her to feel sick, or that it was the combination of keys she pressed that caused the laptop's screen to turn dark, or that it was the storm last night that caused the flower pot to be shattered into pieces? To put it more generally, how do reasoners assess whether a singular instantiation of a cause factor C was actually causally connected to a singular instantiation of its effect E ?

Judgments about singular causation are so prevalent in everyday life that it is easily missed that validating them is a challenging task for the mind. The reason why it is cognitively challenging is that causal powers that bind some events together are not directly accessible to our senses (Cartwright, 1989; Cheng, 1997). The computational problem that needs to be solved is how a genuine *causal co-occurrence* of events can be discriminated from a mere *coincidental* one.

Stephan and Waldmann (2018) have proposed a computational model intended to provide a solution to this problem. Their model is a generalization of Cheng and Novick's (2005) power PC model of causal attribution, which itself relies on Cheng's (1997) power PC theory. Cheng (1997) has shown how the unobservable powers of causes, operationalized as the probability with which causes generate their effects in the hypothetical absence of alternative causes, can be inferred from observable covariation data. Cheng and Novick (2005) adopted this framework and have applied it to the question of how causal power knowledge ought to be used to make

causal attributions in different contexts. For example, when it is known that a potential cause C and a potential effect E have co-occurred on an occasion (c, e) , their model provides an answer to $P(c \rightarrow e|c, e)$, the probability that c and e were causally connected on this occasion.

Stephan and Waldmann (2018) criticized and refined Cheng and Novick's (2005) model. Cheng and Novick's model focuses solely on the causal powers of the target and the alternative causes, embodying the assumption that a singular instantiation of C and E is less likely to be coincidental when C is known to operate with a large causal power. However, the model neglects another possibility why a singular co-occurrence of C and E might have been coincidental: causes, no matter how strong their power is, can be *preempted* in their efficacy by competing alternative causes (for an overview of theories on preemption, see, e.g., Paul & Hall, 2013). It is possible that an alternative cause intercepts the target cause and generates the effect before the target cause has had a chance. Stephan and Waldmann (2018) have therefore proposed a refined model that includes a term that captures the probability of preemption through alternative causes.

To estimate the probability that a target cause was preempted, what needs to be considered beyond the powers of the potential causes is temporal information. One type of temporal information that should influence how strongly a reasoner believes that a target cause was preempted by an alternative cause is the assumed difference between their instantiation times: everything else being equal, a target cause is more likely to be preempted by an alternative cause if the latter occurs earlier than the former. Stephan and Waldmann (2018) reported a set of experiments in which the cover stories suggested that unobserved alternative causes occurred prior to the target cause, and participants' singular causation judgments were explained well by the modified model incorporating the possibility of preemption.

Another type of temporal information that is likewise relevant to assess the probability of preemption is information about *causal latency*, by which we mean the time it takes a cause to produce its effect. Consider a situation in which a potential alternative cause A is instantiated simultaneously with or even later than the target cause C . In such situations, c can still be preempted by a when a 's latency is shorter than c 's.

Stephan and Waldmann (2018) did not spell out in their article how causal latency information about the competing potential causes of an outcome can be formally represented and combined with causal power information to es-

timate $P(c \rightarrow e|c, e)$ and the probability of preemption. Nor did they manipulate causal latency information in their experiments. We begin to address these shortcomings in the present research, and will primarily focus on the role causal latency plays for singular causation judgments. First we will review Cheng and Novick's (2005) model and contrast it with Stephan and Waldmann's (2018) modified version.

The Power PC Model of Causal Attribution

Cheng and Novick (2005) proposed a model in which the probability with which an instantiation c of a cause factor C has actually caused a token event e instantiating an effect factor E is represented by $P(c \rightarrow e|c, e)$. To estimate this probability, they have made use of Cheng's (1997) causal power theory.

Cheng (1997) has shown how the power of a cause factor C , denoted by w_C , can be estimated from observable covariation data given a number of causal background assumptions. A graphical representation of the unobservable causal structure that is assumed by causal power theory to underlie the observable contingency ΔP between a target cause C and a target effect E is shown in Fig. 1 (see also Griffiths & Tenenbaum, 2005; Pearl, 2000). The theory considers two causal influences on E : the target cause C , and A , with A representing the sum of all unobserved alternative causes of E . It is assumed that C and A occur independently of each other with the base rates b_C and b_A . Furthermore, C and A are assumed to cause E with independent (i.e., non-interacting) powers, denoted by w_C and w_A , respectively. These assumptions allow it to explain the probability of E : $P(E) = b_C \cdot w_C + b_A \cdot w_A - b_C \cdot w_C \cdot b_A \cdot w_A$. Accordingly, the probability of E given C is $P(E|C) = w_C + b_A \cdot w_A - w_C \cdot b_A \cdot w_A$, and the base rate of E is $P(E|\neg C) = b_A \cdot w_A$. The latter two equations provide a causal explanation of the observed contingency: $\Delta P = w_C + b_A \cdot w_A - w_C \cdot b_A \cdot w_A - b_A \cdot w_A$. Substituting $b_A \cdot w_A$ with $P(E|\neg C)$ and re-arranging the equation, one obtains the causal power of C :

$$w_C = \frac{\Delta P}{1 - P(E|\neg C)}. \quad (1)$$

As an illustration, consider the table in Fig. 1 with the following entries: $n(c, e) = 21$, $n(c, \neg e) = 3$, $n(\neg c, e) = 12$, and $n(\neg c, \neg e) = 12$. Imagine these data resulted from a study testing whether a drug causes nausea as a side effect. In the control group ($\neg c$), 12 of 24 subjects developed nausea. The causal power theory assumes that these cases are due to the unobserved factors included in A . In the treatment group 21 subjects had nausea. As A is supposed to occur equally likely in C 's presence, these cases are explained by the joint influence of C and A . Based on the independence assumption we can infer that 12 of the 21 subjects with nausea would have developed nausea due to A alone had C not been present. Thus, there remained 9 subjects in which C had the chance to reveal its power. As there are 21 subjects with nausea, we can therefore conclude that the drug *exclusively* caused nausea in nine of these 12 cases. Its causal power thus is $w_C = .75$.




Figure 1: The relation between observable covariation data and unobservable causal structure. C and E in the causal structure denote the cause and effect factor, respectively. A comprises all unobserved alternative causes of E . b_C and b_A denote the base rates of C and A ; w_C and w_A denote the causal powers of C and A .

Now imagine you learned that a person took the drug (c) and felt sick (e). To compute $P(c \rightarrow e|c, e)$, Cheng and Novick (2005) proposed the following equation:

$$P(c \rightarrow e|c, e) = \frac{w_C}{w_C + w_A - w_C \cdot b_A \cdot w_A} = \frac{w_C}{P(E|C)}. \quad (2)$$

In this equation the power of C is in the numerator and the conditional probability of E given C is in the denominator. Stephan and Waldmann (2018) argued that what this equation thus estimates is the relative frequency of cases among all observed co-occurrences of C and E in which C 's power was probabilistically sufficient for E . To see this, consider first how often E occurred in the presence of C in our example. This was the case in 21 of the 24 test group subjects. Hence, we obtain $P(E|C) = .875$. From Eq. 1 we know that C has a power of $w_C = 0.75$, implying that C was sufficient in 75 percent of the treatment subjects (in 18 of the 24). Applying Eq. 2, we see that model thus concludes that C caused E in 18 of the 21 cases in which C and E co-occurred, yielding $P(c \rightarrow e|c, e) = \frac{18}{21} = 0.86$. The probability of a singular co-occurrence having been causal is supposed to be 86 percent.

A Modified Model Sensitive to Preemption

Stephan and Waldmann (2018) have argued that the power PC model of causal attribution tends to overestimate the probability of singular causation. Their argument was that not on every occasion on which a cause factor C is probabilistically sufficient to generate E it has actually caused E because it is possible that C has been preempted by an alternative cause factor A that succeeded in causing E before C could have taken effect. To illustrate the problem, imagine the extreme case in which all alternative factors A generated their effects before C . In this case we would say that there are 12 subjects with nausea due to A in each group, and it seems natural to say in this situation that C actually caused nausea only in those nine subjects that were added to these 12. The probability that a subject from the test group suffered from nausea due to the drug would hence be $P(c \rightarrow e|c, e) = \frac{9}{21} = 0.43$, and not 0.86. The power PC model of causal attribution seems to yield the correct result only in a situation in which C produces all its effects prior to A , or when the two causes are in a relation of *symmetric* overdetermination.

To incorporate the possibility of preemption by alternative causes, Stephan and Waldmann (2018) have proposed the following refined equation:




Figure 2: Pairs of gamma distributions contrasted in the five conditions of Exp. 1. The shape (κ) and scale (θ) parameters of the five different gamma distributions are listed for each pair. The depicted α values were obtained using the MC algorithm corresponding to Eq. 4. The dark distributions show the causal latencies of the target cause and the light distributions show the causal latencies of the alternative cause.

$$P(c \rightarrow e|c, e) = \frac{w_C - w_C \cdot b_A \cdot w_A \cdot \alpha}{P(E|C)} = \frac{w_C \cdot (1 - b_A \cdot w_A \cdot \alpha)}{P(E|C)}. \quad (3)$$

Eq. 3 extends the numerator of Eq. 2 by subtracting from w_C the product of w_C , b_A , w_A , and a newly introduced parameter α . The product of b_A , w_A , and α captures the probability of preemption. The intersection of w_C , b_A , and w_A identifies the occasions in which C and A are both probabilistically sufficient to generate E . This part of the term is relevant because the problem of preemption occurs only on occasions in which the potential causes C and A are both probabilistically sufficient to generate the effect. On those occasions it can either be the case that C preempts A , that A preempts C , or that both act synchronously (discussed in the philosophical literature under the term *symmetric overdetermination*, see, e.g., Paul and Hall 2013). Stephan and Waldmann (2018) introduced the α parameter as a weighting factor that narrows down the intersection of w_C , b_A , and w_A to those occasions on which C is preempted by A . To illustrate the idea, consider again the nausea example. In this example, $w_C \cdot b_A \cdot w_A = 0.75 \cdot 0.50 = 0.375$. Now imagine again the extreme case in which the factors in A produced their effects prior to C . In this scenario, all 37.5 percent of the cases in which C and A were both sufficient for E were actually caused by A . Such a situation can be modeled by setting α to 1. This would yield $P(c \rightarrow e|c, e) = \frac{0.75 - 0.375 \cdot 1.0}{0.875} = 0.43$. By contrast, in the extreme cases in which C generates its effects prior to A or synchronously with A , α should be set to 0. In this case Eq. 3 reduces to Eq. 2.

Incorporating Causal Latency Information

We showed in the previous sections how the possibility of preemption can be incorporated into a model of causal attribution. The open question is how α can be estimated in different contexts. This is where temporal information about the potential causes comes into play.

We will focus on the situation in which A only consists of a single alternative cause factor of the effect. As pointed out above, an obvious factor influencing α in such situations is the difference between the onset times of C and A . This difference can be represented by $\Delta_t = t_a - t_c$. Everything else being equal, when A occurs earlier than C , it is more likely that A preempts C than vice versa.

Additionally, the probability of preemption is influenced by the causal latency of the potential causes. By causal latency we mean the time it takes a cause to produce its effects. Variation in the latency with which a cause generates its effect opens up the possibility that even when C and A are instantiated simultaneously, or when A is instantiated later than C , C could still be preempted by A .

To model causal latencies we will use gamma distributions, which are, for example, used in queueing theory to model waiting times. In recent studies, Bramley, Gerstenberg, Mayrhofer, and Lagnado (in press) used gamma distributions to model the role of time in causal structure induction (see also Lagnado & Speekenbrink, 2010). A gamma distribution is a continuous probability distribution 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 pairs of gamma distributions that we contrasted in our experiments are depicted in Fig. 2.

The representation of causal latencies based on gamma distributions can be used to estimate α in different types of situations. We will here focus only on situations in which it is known that C , A , and E are all present. Moreover, the causal latency distributions of C and A and the size of Δ_t are known. In these situations α corresponds to:

$$\alpha = P(t_{A \rightarrow E} + \Delta_t < t_{C \rightarrow E} | e, c, a, \Delta_t). \quad (4)$$

In this equation, $t_{C \rightarrow E}$ and $t_{A \rightarrow E}$ denote the causal latencies of C and A (cf. Bramley et al., in press), which are given by the respective gamma distributions. In this situation α corresponds to the probability that the sum of the causal latency of the competing cause A and the time lag between C 's and A 's onset times is smaller than the causal latency of C , given e , c , a , and Δ_t . This probability can be estimated with the following Monte Carlo (MC) algorithm:

1. Sample N pairs of causal latencies (x_c, x_a) from the gamma distributions of C and A , respectively.
2. Calculate $x_{a'} = x_a + t_a - t_c$ for all sampled x_a -values.
3. Count all pairs for which $x_c > x_{a'}$.
4. Divide this count by N .

The values for α depicted in Fig. 2 were obtained by applying this algorithm with $N = 10,000$.

Experiment 1

The goal of Exp. 1 was to compare different scenarios intended to manipulate α and to compare the predictions of our model for these situations with people’s singular causation judgments. To simplify the task, and to isolate the influence of causal latency, we used a test scenario in which the competing cause factors occurred simultaneously (i.e., $\Delta_t = 0$). Furthermore, we considered only deterministic causes in this first experiment. Fig. 2 shows the five gamma distributions (G1 - G5; higher numbers indicate higher expected values) that we contrasted in five conditions. The latency distributions belonging to the target cause C in each condition are depicted in dark blue. For the first pair, for example, in which we contrasted G1 and G5, the target cause factor’s latency follows G1, whereas the latency of the alternative cause follows G5. Fig. 2 also shows the different α values estimated with the algorithm presented above. For the first pair $\alpha = 0.01$, which in the case of deterministic causes directly corresponds to the probability that C was preempted by A . Thus, participants should be confident in this condition that it was indeed the target cause C that brought about the observed outcome. In the fifth condition, by contrast, in which C followed G5 and A followed G1, participants should be confident that C did not cause the effect. Fig. 2 also shows that all the other conditions should elicit more uncertainty. In the third condition, for example, in which C and A have the same latency distribution (G3), $\alpha = 0.50$. Here participants should be maximally uncertain about the singular cause of the outcome.

Fig. 4 shows the predictions of the power PC model of causal attribution (Eq. 2) as well as those of our refined model that computes α according to the algorithm corresponding to Eq. 4. We labeled this model “Alpha Precise” because we also considered an “Alpha Coarse” model. Alpha Coarse estimates only whether the expected values of the competing distributions differ and neglects their variances. It thus assigns a value of 1 to α if the expected value of the causal latency distribution of the target cause is larger than the one of the competing cause, and 0 otherwise. Alpha Coarse treats the situation in which both causes follow identical latency distributions as a case of *symmetric overdetermination*, and therefore predicts that both cause factors should be seen as singular causes in this situation.

Participants

Two hundred subjects ($M_{age} = 27.14$, $SD_{age} = 8.71$, 119 females) who had at least an A-level degree and who were native English speakers were recruited via Prolific (www.prolific.ac). Subjects were paid £0.80 for their participation.

Design, Materials, and Procedure

Participants were randomly assigned ($n = 40$) to one of five key conditions. These conditions varied with respect to the contrasted gamma distributions and with respect to the gamma distribution associated with the target cause (see Fig. 2). Because we included several balancing factors, the




Figure 3: Illustration of the learning task used in the experiments.

full design was a $2 \times 2 \times 2 \times 5$ between-subjects design. The additional balancing factors will be introduced below.

Subjects were presented with a scenario about a fictitious medieval kingdom called “Extonia”. They read that the king had two watchtowers (“North” and “South”) built at the border to protect his empire from barbarians. These towers were instructed to send carrier pigeons to the palace to cause alarm whenever barbarians are spotted. Participants were then asked to take the perspective of Extonia’s secretary of defense who routinely inspects the flight durations of the pigeons from the two towers. Participants read that the flight durations tend to differ between different pigeons, and that they will therefore observe a sample of thirteen pigeons from each tower. Before participants could proceed to the learning task, they had to pass an instruction check.

During the learning task the screen looked similar to the picture in Fig. 3. Whether participants began with tower “North” or “South” was balanced between subjects. In each trial the sending of a carrier pigeon was indicated with a delay of 500ms by a circling of the watchtower. The arrival of the pigeon was indicated by a colored circle surrounding the palace. The thirteen flight durations for each tower corresponded to thirteen quantiles of the respective gamma distribution. We used quantiles so that subjects could be presented small but yet representative samples. After each trial, participants had to click a “Next” button, which was operational 500ms after the circle around the palace had been displayed. The flight durations were presented in random order.

The test scenario described a singular situation in which the palace had been alarmed by a tower so that it was possible to repel a horde of barbarians. Participants read that the people of Extonia wanted to decorate the tower that was responsible for the alarm but that there was the problem that both towers had actually sent their pigeons simultaneously. To express their opinion about which tower was the actual cause of the alarm, participants indicated on an eleven-point rating scale (end points: “Definitely not caused by Tower ‘North/South’” and “Definitely caused by Tower ‘North/ South’”; midpoint “50:50”) how strongly they believed that the alarm was caused by Tower “North”/ “South”. Whether the target cause was tower “North” or “South” was balanced between subjects. The orientation of the rating scale was also balanced between subjects.




Figure 4: Model predictions and results (mean singular causation ratings and 95% bootstrapped CIs) for the different conditions of Exp. 1. The predictions of the Power PC Model were obtained from Eq. 2. The predictions of the Alpha Coarse Model are based on α being obtained through an ordinal ranking of the expected values of the compared gamma distributions. For the predictions of the Alpha Precise Model α was calculated with the MC method corresponding to Eq. 4.

Results and Discussion

The results are shown in Fig. 4. Participants' singular causation ratings ($M_1 = 0.91$, $SD_1 = 0.10$, $M_2 = 0.69$, $SD_2 = 0.22$, $M_3 = 0.54$, $SD_3 = 0.24$, $M_4 = 0.37$, $SD_4 = 0.19$, $M_5 = 0.11$, $SD_5 = 0.12$, from left to right) followed a negative linear trend. A polynomial trend analysis confirmed the negative linear trend, $F(4, 195) = 111.70$, $p < .001$, $r = .83$. No other polynomial trend was significant.

These results are at odds with the predictions made by the power PC model of causal attribution (Eq. 2). Subjects took into account information about the causal latency of the potential cause factors to derive singular causation judgments. Moreover, the singular causation judgments followed the predictions of the Alpha Precise model, which utilizes gamma distributions to represent causal latencies. The correlation between model predictions and results was high, $r = .99$, and statistically significant $t(3) = 14.88$, $p < .001$.

Experiment 2

In Exp. 1 we tested deterministic causes because we aimed to isolate the influence that causal latency exerts on singular causation judgments. As Eq. 3 shows, however, the probability of preemption is given by the product of the causal powers and α . In Exp. 2 we therefore studied probabilistic causes. We tested a scenario in which both causes either had a causal power of $w_C = w_A = 0.83$ or of $w_C = w_A = 0.5$. Additionally, we manipulated α by using the first pair of gamma distributions (G1 vs. G5) shown in Fig. 2. We tested this combination of causal power and latency because it leads to an interesting interaction effect, depicted in Fig. 5. When both causes have a high causal latency, our model predicts that ratings in the high-power condition should be lower than in the low-power condition. One reason is that the product that is subtracted from w_C in the high-power condition is so large that the numerator becomes smaller ($0.83 - 0.83 \cdot 0.83 \cdot 0.99 = 0.15$) than in the low-power condition ($0.5 - 0.5 \cdot 0.5 \cdot 0.99 = 0.25$). Furthermore, the denominator is smaller in the low-power condition than in the high-power condition ($0.5 + 0.5 - 0.5 \cdot 0.5 = 0.75$ vs. $0.83 + 0.83 - 0.83 \cdot 0.83 = 0.97$), which means that the numerator in the low-power condition is increased more strongly than in the high-power condition. When both causes have a low causal latency (G1), by contrast, the product that is subtracted in the numerators becomes small, and our model predicts that we should see a reversed order of

judgments. Finally, our model also predicts a main effect of causal latency: causes that tend to precede the efficacy of their competitors should receive higher singular causation ratings than causes whose competitors tend to preempt them.

Fig. 5 also shows the predictions of the power PC model (Eq. 2). As this model is blind to causal latency information, it predicts only a main effect of causal power. Our model does not predict a main effect of causal power.

Participants

One hundred and sixty subjects ($M_{age} = 38.14$, $SD_{age} = 12.20$, 116 females) who had at least an A-level degree and who were native English speakers were recruited via Prolific (www.prolific.ac). They were paid £ 1.20 for participation.

Design, Materials, and Procedure

Subjects were randomly assigned to one of four conditions ($n = 40$) that resulted from a 2 (causal power: $w_C = w_A = 0.83$ vs. $w_C = w_A = 0.50$) $\times$ 2 (causal latency: G1 vs. G5) between-subjects design.

The materials and procedure were largely identical to those of Exp. 1, with the following exceptions: first, we added information about the probabilistic causal nature of the towers to the instructions. Participants read that pigeons might get lost on their way to the palace and that it would hence be important to learn the pigeons' arrival rates. Secondly, the learning task was modified such that causal power information could be conveyed. Other than in Exp. 1, we showed subjects 24 pigeons per tower, to ensure that all participants observe a sufficient number of "successful" pigeons to be able to learn the causal latencies. Subjects in the high-power condition observed 20 successful pigeons per tower, whereas subjects in the low-power condition observed 12 successful pigeons per tower. The flight durations corresponded to 12 or 20 percentiles of the respective latency distributions. Whenever a pigeon failed to reach the palace, the words "Pigeon probably lost" were displayed five seconds after the pigeon had been sent out, which corresponded to the 99.9th percentile of G5.

The test questions were identical with the ones in Exp. 1. Additionally, on a separate screen we asked participants to estimate the causal powers of the towers, as we wanted to control for the possibility that subjects' representations of causal power might be influenced by causal latency. For example, the tower "North" participants were asked the following question: "Based on what you have learned: how many out of 10




Figure 5: Model predictions and results (mean singular causation ratings) of Exp. 2. Error bars denote 95% bootstrapped CIs.

letter pigeons sent from Tower ‘North’ would make it to the palace?”. Ratings were provided on an eleven-point scale (0 to 10). The test questions were presented in counterbalanced order.

Results and Discussion

The results are summarized in the right panel of Fig. 5. The singular causation ratings followed the pattern predicted by our model, whereas the power PC model of causal attribution cannot explain the results. When α was high, ratings in the low-causal power condition were higher than those in the high-causal power condition. The reversed pattern was obtained when α was low. A planned contrast testing the predicted interaction was significant, $t(156) = 2.68, p < .01, r = .16$. Furthermore, singular causation ratings were overall higher when α was low. A planned contrast testing this predicted main effect was also significant, $t(156) = 10.47, p < .001, r = .63$. There was no main effect of causal power, $t(156) < 1.00$.

The additional causal power ratings showed that participants’ causal power representations were not distorted by the different causal latencies. In the low-causal power condition, the mean ratings for the fast and slow tower were $M_{fast} = 5.18$ ($SD_{fast} = 1.09$) and $M_{slow} = 5.30$ ($SD_{slow} = 1.16$). In the high-causal power condition, the mean ratings were $M_{fast} = 7.55$ ($SD_{fast} = 1.61$) and $M_{slow} = 7.69$ ($SD_{slow} = 1.97$). A mixed ANOVA with “causal-power query” as within-subject factor yielded a main effect only for causal power, $F(1, 156) = 121.42, p < .001, \eta^2_G = .39$, confirming that subjects in the high-causal power condition gave higher causal power ratings than subjects in the low-causal power condition. Importantly, there was neither a main effect of causal latency, $F(1, 156) < 1$, nor an interaction effect of causal latency and causal power, $F(1, 156) < 1$. We can thus rule out that the results were driven by different causal power representations in the different latency conditions.

General Discussion

To assess singular causation, more than just the causal powers of the potential causes need to be considered. Even when a cause factor is sufficient to generate the effect, it is still not the actual cause of the outcome if it was preempted by a competitor. To assess the probability of preemption, temporal information about the potential causes needs to be considered in combination with information about their power. We have discussed the roles of two relevant types of temporal

information: the difference of the onset times of the competing causes, and the latencies with which they generate their effects. To model the latencies of causes we used gamma distributions, which allows us to estimate the size of the α parameter in our model. The results of two experiments that we presented were explained well by our model, but not by Cheng and Novick’s (2005) power PC model of causal attribution which neglects temporal information.

The type of situation we considered here represents only a subset of possible cases. For example, unlike in our scenario reasoners often experience the latency between the target cause and outcome. Consider, for instance, a situation in which a person takes an aspirin and after twenty minutes gets relief from her headaches. In such cases, the experienced delays need to be used to estimate α . We plan to use dynamic test scenarios in future experiments to test such contexts.

Another noteworthy characteristic of our test scenario was that all potential causes of the effect were actually observed. In most real-world situations, however, reasoners observe only a subset of the potential causes. Other than in our test scenario, reasoners often are confronted with uncertainty concerning the presence of alternative causes. Although we have not considered these situations here, our model can be applied to them, too. For example, in situations in which only the target cause is observed but not alternative causes, information about the temporal distribution of the effect in the absence of the target cause needs to be considered to estimate α . This temporal distribution can be modeled with exponential functions (see, e.g., Bramley et al., in press). We plan to investigate such situations in future studies.

Acknowledgments This work was supported by the 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). We thank Louisa Reins for helpful discussions and assistance with the experiments.

References

- Bramley, N. R., Gerstenberg, T., Mayrhofer, R., & Lagnado, D. A. (in press). The role of time in causal structure learning. *Journal of Experimental Psychology: Learning, Memory & Cognition*.
- Cartwright, N. (1989). *Nature’s capacities and their measurement*. Oxford: Clarendon Press.
- Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, 104, 367–405.
- 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.
- Griffiths, T. L., & Tenenbaum, J. B. (2005). Structure and strength in causal induction. *Cognitive Psychology*, 51, 334–384.
- Lagnado, D. A., & Speekenbrink, M. (2010). The influence of delays in real-time causal learning. *The Open Psychology Journal*, 3, 184–195.
- Paul, L. A., & Hall, E. J. (2013). *Causation: A user’s guide*. Oxford University Press.
- Pearl, J. (2000). *Causality: Models, reasoning and inference*. Cambridge, England: Cambridge University Press.
- Stephan, S., & Waldmann, M. R. (2018). Preemption in singular causation judgments: A computational model. *Topics in Cognitive Science*, 10, 242–257.
- Waldmann, M. R. (Ed.). (2017). *The Oxford handbook of causal reasoning*. New York: Oxford University Press.