

Preemption in Singular Causation Judgments: A Computational Model

Simon Stephan (simon.stephan@psych.uni-goettingen.de)

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

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

Abstract

Causal queries about singular cases are ubiquitous, yet the question of how we assess whether a particular outcome was actually caused by a specific potential cause turns out to be difficult to answer. Relying on the causal power approach, Cheng and Novick (2005) proposed a model of causal attribution intended to help answering this question. We challenge this model, both conceptually and empirically. The central problem of this model is that it treats the presence of sufficient causes as necessarily causal in singular causation, and thus neglects that causes can be *preempted* in their efficacy. Also, the model does not take into account that reasoners incorporate uncertainty about the underlying causal structure and strength of causes when making causal inferences. We propose a new measure of causal attribution and embed it into our structure induction model of singular causation (SISC). Two experiments support the model.

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

Introduction

Most people hold the belief that smoking causes lung cancer. Now, imagine that you learn that Peter, a passionate smoker, has contracted lung cancer. How strongly would you be willing to say that it was Peter's smoking that was causally responsible for his disease?

This example illustrates a scenario in which we seek an answer to a causal query about a singular case. Queries about singular causation are prevalent in everyday life and professional contexts, such as the law or medicine. How do people derive causal judgments about singular cases? Of course, the mere fact that two factors C and E are *generally* causally connected (e.g., smoking often causing lung cancer) does not necessarily imply that a *singular* or *token* co-occurrence of these events (e.g., Peter's smoking and his lung cancer) manifests a causal relationship – a singular co-occurrence might be a mere *coincidence*. On the other hand, as causality is not directly observable in the world, to what else than our general causal knowledge could we turn to obtain answers?

We are going to present a theory that builds on the idea, first formalized by Cheng and Novick (2005), that the notion of unobservable *causal powers* (Cheng, 1997) plays an essential role in singular causation judgments. Yet, we will demonstrate that Cheng and Novick's (2005) power PC model of causal attribution (CN model) makes assumptions that are not always plausible. Generally, the CN model is intended to provide a normative answer to the question how we can determine whether an observed outcome was actually caused by a potential cause factor. For example, for cases like the one above about Peter in which a potential cause c and an effect e have been observed, the CN model delivers the probability

$P(c \rightarrow e|c, e)$ with the arrow denoting a causal relation. We argue that the key problem of the model is that it treats target causes as singular causes whenever they are sufficient for the effect in a specific situation. This appears to be at first sight a reasonable assumption, yet it ignores that the exact points in time at which different causes exert their powers play an important role in singular causation judgments (see Danks, 2017): crucially, sufficient causal powers can be *preempted* by others, and in such cases they should not be held causally responsible for the occurrence of the outcome. We will argue that preemption of causes by background factors frequently occurs in singular causation scenarios, and therefore presents a problem for the CN model.

Another problem of the CN model is that it does not take into account uncertainty about both the underlying causal structure and the causal parameters (e.g., the size of the causal powers). To incorporate uncertainty about the causal parameters, Holyoak, Lee, and Lu (2010) have proposed a Bayesian version of the CN model that uses probability distributions over the parameters instead of point estimates. However, their model also neglects uncertainty about the underlying causal structure. Both sources of uncertainty have been demonstrated to influence causal learning and reasoning (Griffiths & Tenenbaum, 2005; Meder, Mayrhofer, & Waldmann, 2014). For this reason, Stephan and Waldmann (2016) proposed the structure induction model of singular causation (SISC) that incorporates both types of uncertainty. Although three experiments (Stephan & Waldmann, 2016) showed that SISC better accounted for the results than the standard power PC model of causal attribution, one shortcoming of the initial version of SISC was that it used the CN conceptualization of causal attribution that we are going to criticize in the present paper.

We will start with a theoretical section in which we defend a new measure of causal attribution as a component of SISC that is sensitive to preemption. We then present the results of two experiments. Experiment 1a confirmed that singular causation judgments deviate systematically from the predictions of the CN model in line with our revised causal attribution equation. Experiment 1b assessed participants' notion of preemption. In Experiment 2 we used a larger set of contingencies to compare the revised SISC with the CN and other models. The results of this experiment showed that both a revision of the causal attribution equation and the consideration of statistical uncertainty are crucial to explain the findings.

The Power PC Model of Causal Attribution

According to Cheng's (1997) power PC theory, causal power (or causal strength) is a hypothetical, unobservable en-

tity that represents the strength of causes. Mathematically, causal power is (in the generative case) expressed as the probability with which a target cause brings about its effect in a hypothetical world in which all alternative observed and unobserved causes of the effect are absent. Because of the possibility of unobserved alternative causes, causal power cannot be assessed directly but must be inferred based on the observed covariation and background assumptions. For generative causes, the following equation can be used to estimate the causal power w_c of a target cause C :

$$w_c = \frac{P(e|c) - P(e|\neg c)}{1 - P(e|\neg c)} = \frac{\Delta P}{1 - w_a} \quad (1)$$

In this equation, w_a represents the aggregate causal power of all alternative causes A of the effect, which are assumed to exert their influence independently of C .

Under the causal Bayes net framework, the causal power of C , w_c , corresponds to the probabilistic weight of the causal arrow that connects C with its effect E in a common effect structure in which the target cause C and the alternative causes A combine in a *noisy-OR* gate (see S_1 in Figure 2). Likewise, w_a corresponds to the weight of node A .

The CN Measure of Causal Attribution

Cheng and Novick (2005) proposed several measures of causal attribution that apply to different cases. The measure of causal attribution for cases in which both c and e are present, as in the example above about Peter, utilizes the concept of causal power in the following way to deliver the conditional probability $P(c \rightarrow e|c, e)$:

$$P(c \rightarrow e|c, e) = \frac{w_c}{P(e|c)} = \frac{w_c}{w_c + w_a - w_c \cdot w_a}. \quad (2)$$

Equation 2 shows that the CN model defines the probability with which c is causally responsible for e given that both have co-occurred by the fraction of the causal power of C and the conditional probability of the effect in the presence of C . Since the power PC theory assumes that C and A exert their causal powers independently of each other, $P(e|c)$ can be rewritten as the sum of both causal powers minus their intersection (see second step of Equation 2). Hence, what the CN model delivers is an estimation of the relative frequency of cases among all co-occurrences of C and E in which C 's causal power is sufficient for the production of the effect. Our key criticism is that this relative frequency, because it neglects the possibility of preemption, frequently overestimates the true proportion of cases in which we should actually causally attribute E 's occurrence to C .

To illustrate the problem, let us consider the results of the fictitious experiment shown in Figure 1 in which the influence of a chemical substance on the expression of a gene was investigated. As it is the case that all mice in the test group ($P[e|c] = 1$) but only one half in the control group ($P[e|\neg c] = .5$, the base rate) exhibit the gene, the results provide strong evidence for the existence of a strong effect of the chemical. In fact, by applying Equation 1 one can see that

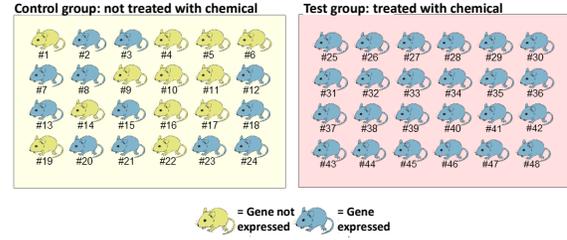


Figure 1: Illustration of a hypothetical study testing the effect of a chemical on the expression of a gene. The control group is shown on the left and the test group treated with the chemical substance on the right. Mice having the gene expressed are depicted in blue.

the causal power of the substance equals 1. Crucially, so does $P(c \rightarrow e|c, e)$. What the CN model therefore prescribes is that we should attribute causal responsibility to the chemical whenever the chemical and the gene are both present. But should we really be maximally confident that the expression of the gene in, for instance, Mouse #25 must be causally attributed to the causal power of the chemical? If you have doubts, you were probably led by a prior assumption about the point in space and time at which the causal background factors A produced the observed base rate of fifty percent: in the given scenario it seems likely that these factors (e.g., transcription factors) already produced their effects prior to the introduction of the chemical. Under this assumption, however, it seems likely that not only fifty percent of the mice in the control condition but also in the test group already possessed the gene prior to the study. Consequently, in those fifty percent of the mice it cannot be the chemical that is causally responsible for the effect because its causal efficacy has been preempted by the background factors.

A New Measure of Causal Attribution

In the example it seems appropriate to say that the expression of the gene is caused by the chemical in only about half of the observed cases in which C and E have co-occurred. This conclusion is based on the assumption that in roughly half of the cases the causal power of the chemical has been preempted by the causal power of the background factors A . We propose a new measure of causal attribution that captures this intuition by refining Equation 2 so that all cases among the joint occurrences of C and E for which the effect of C is assumed to be preempted by A are partialled out. This refined measure is given by:

$$P(c \xrightarrow{\text{sing.}} e|c, e) = \frac{w_c \cdot (1 - \alpha \cdot w_a)}{w_c + w_a - w_c \cdot w_a}, \quad (3)$$

in which we introduce α as a discounting parameter that represents the assumed probability with which A is a preemptive cause of the effect. For illustration, if we assume that A has caused the observed base rate of 0.5 prior to the application of the chemical, and if we assume further that A 's causal power has produced roughly equal proportions of the effect in both groups, α takes on a value of 1. In this case, the point estimate for $P(c \xrightarrow{\text{sing.}} e|c, e)$ is about .5 instead of 1.0 that is predicted

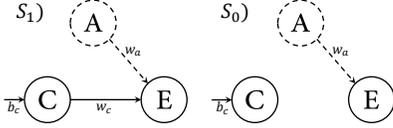


Figure 2: The two causal structures considered by the structure induction model of singular causation as mutually exclusive explanations for observed patterns of covariation. C is a general cause of E under S_1 whereas all co-occurrences of C and E are coincidental under S_0 . The parameters w_c and w_a denote the causal powers of the observed cause C and the unobserved background causes A , respectively; b_c denotes the base rate of C .

by Equation 2. The difference between the two measures is that under the CN model the presence of two sufficient causal powers (w_c and w_a) is invariably conceptualized as a case of “symmetric overdetermination”, whereas the possibility that causes can preempt each other is neglected. Equation 3 takes into account the possibility of preemption and thus delivers an estimation of the relative frequency of singular cases in which the target cause has *actually* been successful in generating the effect. In our view, preemption of C by a previously present background factor A seems to be prevalent in the cover stories typically reported in the literature (see e.g., Griffiths & Tenenbaum, 2005). However, the discounting parameter α can be set to capture other cases. For example, cases of overdetermination or cases in which A is preempted by C could be modeled by setting α to 0. In these cases, our model and the CN model make identical predictions because Equation 3 reduces to Equation 2.

SISC: The Structure Induction Model of Singular Causation

Apart from the fact that the CN model does not take into account the possibility that preempted causes should not be classified as singular causes, a further problem of the CN model is that it is insensitive to statistical uncertainty about both the underlying causal structure and the size of the causal parameters. SISC (Stephan & Waldmann, 2016) is sensitive to both types of uncertainty.

SISC was developed in the framework of causal Bayesian inference models; it takes observed data as evidence to update prior probabilities of mutually exclusive hypotheses. Under SISC, these competing hypotheses represent two causal structures that can account for a particular observed pattern of covariation. The two causal structures, S_0 and S_1 , are depicted graphically in Figure 2. While there exists a causal arrow from C to E in S_1 , which indicates that C is a general cause of E , there is no causal arrow between C and E in S_0 . Both models assume a background cause A .

The core principle of SISC can be illustrated with Figure 1. Assume someone suggests that S_0 is the causal structure that underlies the results. Under this hypothesis all observed co-occurrences of C and E would be mere coincidences. Yet, since the observed distribution of the events appears very unlikely to be coincidental, S_0 is weakened as an explanation while the alternative hypothesis, S_1 , is proportionally strengthened. In fact, the probability computed by SISC for

S_1 for the data shown in Figure 1 (i.e., the posterior probability of S_1) is almost 1. Now, imagine the same study had been conducted with a sample of merely eight mice but that $P(e|c)$ and $P(e|\neg c)$ remain the same. In this case, it seems less certain that S_1 underlies these results. Smaller samples not only increase uncertainty about the underlying causal structure, they also impede the reliable estimation of the size of parameters. SISC is sensitive to both types of uncertainty when estimating $P(c \xrightarrow{\text{sing.}} e|c, e)$.

SISC implements different steps. First, it derives the posterior probabilities for each causal structure illustrated in Figure 2. Applying Bayes’ rule, the posterior probability for a causal structure is proportional to the likelihood of the data given the causal structure, weighted by the structure’s prior probability:

$$P(S_i|D) \propto P(D|S_i) \cdot P(S_i). \quad (4)$$

$P(D|S_i)$ is the likelihood of the data given a particular structure, which is the integral over the likelihood function of the parameter values under the particular structure. $P(S_i)$ represents a structure’s prior probability. The model initially assumes that both structures are equally likely, that is, $P(S_i) = 1/2$. When data become available, the posterior for each causal structure varies systematically with the observed contingency: the higher the contingency, the more likely S_1 becomes.

Next, the model estimates the parameters b_c , w_c , and w_a , for each causal structure. To express parameter uncertainty, distributions rather than point estimates are inferred. The posterior probability distributions for the parameters, $P(w|D)$, are proportional to the likelihood of the data given the set of parameters w , weighted by the prior probability distributions of the parameters:

$$P(w|D) \propto P(D|w) \cdot P(w). \quad (5)$$

$P(D|w)$ is the likelihood of the data given the parameter values for b_c , w_c , and w_a . $P(w)$ is the prior joint probability of the parameters. The prior distributions of the parameters are independently set to flat, uninformative $beta(1,1)$ distributions. Since C does not cause E under S_0 , w_c is held fixed at 0 for this causal structure.

In the last step, SISC computes $P(c \xrightarrow{\text{sing.}} e|c, e)$ for each parameterized structure. The new discounting parameter alpha is set based on background assumptions about the target scenario. For the scenarios we used in the present experiments it is set to 1 because preemption seems to be highly probable. As all co-occurrences of c and e are coincidences under S_0 , $P(c \xrightarrow{\text{sing.}} e|c, e)$ is set to 0 for S_0 . For S_1 , Equation 3 is applied. The final output of SISC is a single estimate for $P(c \xrightarrow{\text{sing.}} e|c, e)$, which is obtained through integrating out the two causal structures by summing over the derived values of $P(c \xrightarrow{\text{sing.}} e|c, e)$ for each structure weighted by its posterior probability:

$$P(c \xrightarrow{\text{sing.}} e|c, e; D) = \sum_i P(c \xrightarrow{\text{sing.}} e|c, e; S_i) \cdot P(S_i|D). \quad (6)$$

Experiment 1a

The goal of Experiment 1a was to test SISC against the CN model of causal attribution for data sets with a sufficient cause, i.e., $P(e|c) = w_c = 1$, but varying base rates of the effect. Whereas the CN model predicts maximal confidence in singular causation assessments for any observed co-occurrence of C and E in this case, SISC predicts an interaction with the base rate under the assumption that A 's causal power generally preempts the effect of C . The goal of Experiment 1a was to demonstrate that this predicted deviation from the CN model is expected for the conceptual reasons discussed above. To rule out uncertainty as an explanation, we used sample sizes in our data sets for which the posterior probabilities of S_1 computed by SISC are close to 1. The predictions of the models are shown in Figure 3. We set α in Equation 3 to 1, which represents complete preemption of C by A . We also considered a Bayesian variant of the CN model that has been proposed by Holyoak et al. (2010). This model is sensitive to parameter uncertainty; it uses probability distributions over the parameters instead of point estimates. As Figure 3 shows, the predictions of both variants of the CN model converge for large sample sizes because the influence of parameter uncertainty decreases.

Methods

Participants 90 participants ($M_{age} = 33.24$, $SD_{age} = 12.50$, 35 female) were recruited via Prolific Academic (www.prolific.ac) and received a monetary compensation of £0.60.

Design, Materials, and Procedure Three contingencies (see Figure 3) were manipulated between subjects with each participant responding to two causal test queries (general causation vs. singular causation). We included the general causation query to establish that uncertainty cannot account for the predicted pattern of singular causation ratings. The task was a standard elemental causal induction task. As cover story we used the gene expression scenario (cf. Griffiths & Tenenbaum, 2005) mentioned above: subjects were asked to assume that they were biologists who are interested in whether a particular chemical causes the expression of a particular gene in mice. Subjects read that they will be asked to conduct an experiment on the computer screen in which they will treat a random sample of mice with the substance while a control sample will remain untreated. It was mentioned that the control sample is important as some individuals may show the gene expression for other reasons.

Participants were presented with an interactive animation showing the two samples arranged as in Figure 1, and a pipette containing a reddish chemical substance. All mice had gray color in this animation. Participants then dropped the substance into the test group area, whereupon the background color changed to a light red. On the next screen, subjects checked the results of the experimental manipulation by dragging a small magnifying glass over all the mice. Mice with the gene then became blue and those without became yellow. The final state of the animation looked like Figure 1.

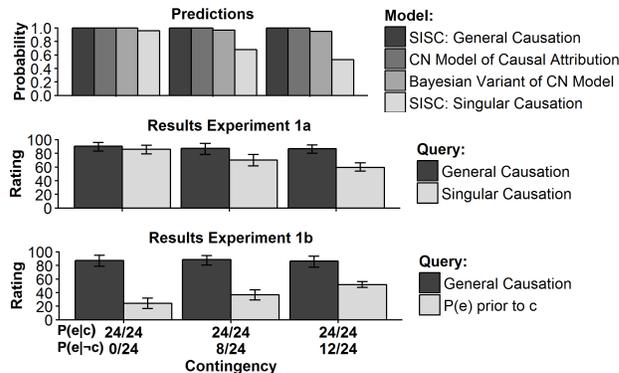


Figure 3: Model predictions and results of Experiments 1a and b. The results show mean ratings and 95% bootstrapped CIs. Dark bars show general causation judgments; light bars singular judgments.

Subsequently, participants responded to two test questions. The general causation query referred to the causal *structure*. Participants were asked to indicate on a slider how confident they were that the chemical has an effect on the expression of the gene (from “very certain that the chemical has no effect” to “very certain that the chemical has an effect”). The singular causation query asked subjects about Mouse #25 from the test group. Participants were asked to indicate on a slider how confident they were that it was the chemical substance that caused the expression of the gene in this single case (from “very certain that it was not the chemical” to “very certain that it was the chemical”).

Results and Discussion

Figure 3 shows the results. The prediction for general causation responses corresponds to the posterior probability of S_1 computed by SISC. As predicted by the posterior probability of S_1 , all general causation ratings were high, indicating very little uncertainty about the general causal structure. The singular causation ratings, by contrast, decreased with an increasing base rate of the effect, as predicted by SISC but not by the two CN models. The results of a multilevel model analysis revealed significant main effects for type of causal query, $\chi^2(1) = 32.45$, $p < .001$, as well as contingency, $\chi^2(1) = 12.63$, $p < .01$. General causation ratings were, on average, higher than singular causation ratings. Figure 3 shows that the main effect of contingency is driven by the decrease in singular causation ratings. Planned contrasts revealed that the general causation ratings neither differed between the first and second contingency, $t(80) = 0.60$, nor between the second and third contingency, $t(80) = 0.13$. Consequently, the interaction effect of query \times contingency was also significant $\chi^2(1) = 13.10$, $p < .01$. Planned contrasts breaking down this interaction effect showed that the difference between general and singular causation ratings was higher for the second than for the first contingency, $t(80) = 2.10$, $p < .05$, $r = .23$, and also higher for the third compared to the second contingency, $t(80) = 3.70$, $p < .001$, $r = .38$. In sum, both the trends for general as well as for singular causation ratings are captured well by SISC.

The observed trend for the singular causation judgments is, however, neither predicted by the CN model using point estimates nor by the Bayesian extension incorporating parameter uncertainty.

Experiment 1b

The goal of Experiment 1b was to assess how likely participants think it is that a particular individual from the test group already exhibited the effect caused by the background factors prior to the occurrence of the cause. Thus, instead of singular causation judgments for a particular individual, we asked subjects to provide a probability judgment. Crucially, responses to this query provide us with an estimate of the α value in Equation 3 that participants assumed.

Methods

Participants 88 participants ($M_{age} = 31.22$, $SD_{age} = 10.84$, 42 female) participated in this only study and received a monetary compensation of £0.60.

Design, Materials, and Procedure The study design and the materials were the same as in Experiment 1a. The only difference was that, instead of a singular causation judgment for Mouse #25, we asked participants how likely they think it is that this individual already had the gene expressed prior to the experiment. The general causation query remained the same.

Results and Discussion

Figure 3 shows that we replicated the pattern for general causation judgments found in Experiment 1a. Planned contrasts revealed that these ratings did not differ (all t values < 1). However, the probability judgments about the presence of the effect prior to the application of the chemical in the single case showed the opposite trend as the singular causation judgments in Experiment 1a. This finding supports our hypothesis that assumptions about preemption influence singular causation judgments, as predicted by Equation 3. Planned contrasts confirmed that ratings increased from the first to the second, $t(71) = 2.67$, $p < .01$, $r = .30$, and also from the second to the third contingency, $t(71) = 3.16$, $p < .01$, $r = .35$. Furthermore, the results indicate that participants indeed assumed high α values.

Experiment 2

Experiment 1a showed that singular causation ratings for sufficient causes deviate systematically from the predictions of the CN models. This deviation is predicted as a consequence of assumptions about preemption relations between C and A . Experiment 2 pursued two main goals: first, we aimed to test SISC using a larger set of contingencies with a combination of different levels of $P(e|c)$ and $P(e|\neg c)$. Second, we wanted to demonstrate that parameter and structure uncertainty indeed influence general and singular cause judgments. We used the set of contingencies studied in Buehner, Cheng, and Clifford (2003) but excluded the one contingency from the set in which the effect never occurs. It does not make sense to ask for singular causation if the effect is absent. The data sets and model predictions are shown in Figure 4. We set

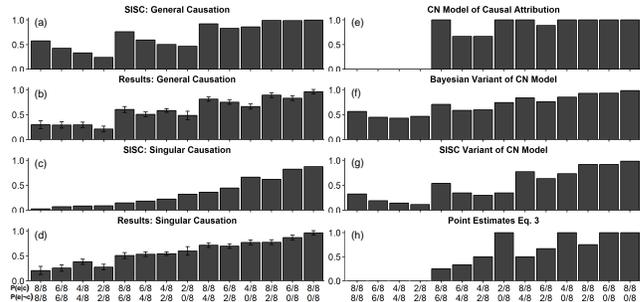


Figure 4: Predictions of different models and results (means and within-subjects adjusted 95% CIs) of Experiment 2. Graphs (a) and (b) refer to general causation assessments. All other graphs refer to singular causation assessments.

the discount parameter α to 1 again.

Methods

Participants 82 participants ($M_{age} = 34.41$, $SD_{age} = 10.42$, 31 female) participated in this online study and were paid £1.00 for their participation.

Design, Materials, and Procedure The causal query (general causation vs. singular causation) was manipulated between subjects, whereas contingency was varied within subject. The fourteen contingency data sets were presented in random order. We used the same cover story as in Experiment 1, except that subjects read that they will investigate the effects of fourteen different chemicals on fourteen different genes in fourteen different samples. We pointed out that the results of the studies are independent of each other. The assignment of mice to the cells of the contingency tables was randomly determined. Also the test mouse for the singular query showing both c and e was randomly chosen prior to the experiment.

Results and Discussion

Figure 4 shows the results and the predictions of the different models: (a) and (b) display the predictions of SISC for general causation and the mean general causation responses. Panels (c) and (d) show the predictions of SISC and the results regarding the singular causation queries. Predictions of the standard CN model and its Bayesian variant are displayed in (e) and (f). Graph (g) shows predictions of SISC when α is set to zero but both structure and parameter uncertainty are incorporated. Finally, (h) shows point estimates of Equation 3 while neglecting statistical uncertainty.

Table 1: Model comparisons for singular causation judgments in Experiment 2. ΔP refers to the different contingency levels (.00, .25, .75, 1.00) within the whole data set; $r_{\Delta P}$ expresses the model fits for these levels. N/A represents undefined values.

Fit measure	SISC	CN Model	Bayesian CN Model	SISC CN Model	Point Est. Eq. 3
$r_{\Delta P=.00}$.72	N/A	-.78	-.68	N/A
$r_{\Delta P=.25}$	1.00	.21	.38	-.61	.98
$r_{\Delta P=.50}$.88	.68	.79	.44	.85
$r_{\Delta P=.75}$	1.00	N/A	1.00	-1.00	1.00
$M_{r\Delta P}$.90	.44	.35	-.46	.94
$r_{overall}$.94	.88	.93	.90	.90
R^2	.88	.77	.87	.82	.82
$RMSE$.11	.27	.09	.14	.22

The overall pattern for both general and singular causation ratings was captured best by the revised version of SISC. As in our previous research, the results show that participants differentiated between general and singular queries. Moreover, the responses to the general causation query replicate those found in Griffiths and Tenenbaum (2005). Most importantly, the singular causation assessments were captured best by the revised SISC. The other models, by contrast, struggled to account for the local trends observed within the subsections of the contingency set in which ΔP is constant. The difference between the singular causation ratings and the point estimates for the revised causal attribution measure (Equation 3) implies that participants were sensitive to structure and parameter uncertainty.

A multilevel model analysis confirmed the main effect of contingency, $\chi^2(13) = 1010.04$, $p < .001$, as well as the interaction between contingency \times query, $\chi^2(13) = 49.39$, $p < .001$, that is shown in Figure 4. To test the different models, we computed different fit measures shown in Table 1. As can be seen there, SISC achieved a good fit in the overall fit measures (bottom part of the table). It explained most variance, with $R^2 = .88$, and yielded the second smallest *RMSE* of .11. Yet, all models obtained relatively high values on the global measures. Even the CN model with the lowest overall fit accounted for 77 percent of the variance. The similarity between the models is not unexpected, however, as all models are sensitive to ΔP . More interesting are the fit measures for the subsections of the contingency set in which ΔP is kept constant. The upper part of Table 1 shows that SISC yielded high fit values there, too, and hence accounted well for these local trends, whereas the Bayesian CN model, which yielded the smallest *RSME*, even showed negative correlations here.

General Discussion

We addressed two different problems that the power PC framework of causal attribution (Cheng & Novick, 2005) faces: first, the CN model attributes causal responsibility for the occurrence of a particular effect e to a present singular event c whenever its causal power is sufficient to bring about the effect. We have argued that this conceptualization fails to take into account that people make assumptions about the point in time at which different causal powers exert their influences. Not every manifestation of a sufficient cause c needs to be causally responsible for an observed outcome; it might be the case that a competing cause (e.g., a) preempts it. This problem of *redundant* causation, which occurs whenever two causes are individually sufficient for the effect, is widely acknowledged in the philosophical literature as a challenge for models of causation (see, e.g., Paul & Hall, 2013). To account for the possibility of preemption we have modified the equation developed by Cheng and Novick (2005) as an account of causal attribution. The revised equation includes the discount parameter α that can be set to express domain-related assumptions about the temporal relations between the alternative causal factors. A second shortcoming of the standard causal attribution model (Cheng & Novick, 2005) is that it

does not take into account statistical uncertainty about structure and causal parameters (cf. Griffiths & Tenenbaum, 2005; Meder et al., 2014). Our model SISC remedies both shortcomings. It is sensitive to both the temporal relations between the alternative causes and to statistical uncertainty. Our experiments showed that both aspects are important to account for subjects' judgments about singular causation.

We have set the discount parameter α to 1 in Equation 3 which implies a complete preemption relation between A and C whenever A 's causal power is sufficient in a situation. Better fits might be possible by estimating the size of α for each individual subject separately. We avoided this strategy to demonstrate that model improvements can already be achieved with very general assumptions. The goal of future experiments will be to manipulate the size of α by manipulating domain assumptions about the temporal relations between C and A . Cases in which α is 1 are situations in which A always preempts C . The cover stories used in the present experiments are an example in which it is plausible to assume that A represents a temporally stable factor that has already been efficacious prior to the manipulation of C . Although preemption seems to be the default situation in most singular causation scenarios, there might be rare cases in which other assumptions need to be made. Consider cases of symmetric overdetermination that have also been discussed in the literature (see Paul & Hall, 2013): in the famous firing squad scenario, for example, in which each shooter is a sufficient cause for the death of the target, a possible intuition is that each shooter should be counted as a singular cause of the death of the victim. In this case, alpha would have to be set to zero. Similarly, alpha would have to be set to zero if C preempts A so that A cannot manifest its potential causal power. Cases of temporal variability between C and A might also be an interesting topic for future studies.

Acknowledgments We thank Jonas Nagel and Ralf Mayrhofer for helpful discussions.

References

- Buehner, M. J., Cheng, P. W., & Clifford, D. (2003). From covariation to causation: A test of the assumption of causal power. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29(6), 1119.
- Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, 104(2), 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.
- Danks, D. (2017). Singular causation. In M. R. Waldmann (Ed.), *The Oxford handbook of causal reasoning* (pp. 201–215). New York: Oxford University Press.
- Griffiths, T. L., & Tenenbaum, J. B. (2005). Structure and strength in causal induction. *Cognitive Psychology*, 51, 334–384.
- 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.
- Meder, B., Mayrhofer, R., & Waldmann, M. R. (2014). Structure induction in diagnostic causal reasoning. *Psychological Review*, 121, 277–301.
- Paul, L. A., & Hall, E. J. (2013). *Causation: A user's guide*. Oxford University Press.
- Stephan, S., & Waldmann, M. R. (2016). Answering causal queries about singular cases. In A. Papafragou, D. Grodner, D. Mirman, & J. C. Trueswell (Eds.), *Proceedings of the 38th Annual Conference of the Cognitive Science Society* (pp. 2795–2801). Austin, TX: Cognitive Science Society.