When do we think that X caused Y?

Tadeg Quillien

October 4, 2019

Contents

1	Portable causal dependence	5
2	A formal model of actual causal strength	9
2.1	Structural equation models	9
2.2	Model	10
3	Testing the empirical fit of the model	13
3.1	Results	15
3.1.1	Conjunctive structure	15
3.1.2	Disjunctive structure	18
3.2	Discussion	20
4	General discussion	21
4.1	Non-statistical normality	22
4.2	Non-monotonic causal effects	23
4.3	The importance of what actually happened	23
4.4	The function of actual causation	26

5	Conclusion	28
	References	29
6	Appendix	32
6.1	Proof that the causal metric reduces to a correlation coefficient under the no-confounding assumption	32
6.2	Analytical expressions for actual causal strength in simple con- junctive and disjunctive causal structures	34
6.2.1	Conjunctive causal structure	35
6.2.2	Disjunctive causal structure	36
6.3	Non-linear patterns in the disjunctive causal structure.	37

Abstract. When explaining an event, people tend to select a single cause out of the multiple factors that contributed – for instance, they will say that a forest fire was caused by a lit match, without mentioning the oxygen in the air which helped fuel the fire. Recently scholars have suggested that causal selection is designed to provide explanations that are likely to generalize across a variety of background circumstances. Here, we develop a computational model of causal selection which formalizes this idea. Under minimal assumptions, the model is surprisingly simple: a factor is regarded as a cause of an outcome to the extent that it is, across counterfactual worlds, correlated with that outcome. The model explains why causal selection is influenced by the normality of candidate causes, and outperforms other known computational models when tested against a fine-grained dataset of human graded causal judgments.

People spontaneously and effortlessly assign causes to specific events; in the jargon of psychologists and philosophers, they easily make judgments of ‘actual causality’. Yet how they do so is still poorly understood. In particular, when multiple causes contribute to an event, we tend to discriminate among them: for instance, we tend to say that the forest fire was caused by the match lit by a careless camper, but we regard the presence of oxygen in the air as a mere ‘enabling condition’ or ‘contributing factor’. This suggests that we implicitly rank the different causes of an event, as if we attributed a ‘causal score’ to each of them [1, 2].

This paper proposes a model of how the mind assigns these causal scores. The model is based on the intuition that actual causation judgments are aimed at providing explanations that generalize beyond the narrow specific circumstances of a particular event [3, 4]. It has two desirable properties.

First, under minimal assumptions, the model is surprisingly simple: an

event C is judged to have caused an event E to the extent that C and E are, across counterfactual worlds, correlated with each other. For instance, people tend to judge a victory in a swing state like Florida as highly causal to the victory of a US presidential candidate. Under our model, this is because when we imagine many alternative ways that the election could have unfolded, the counterfactuals where the candidate loses the presidency also mostly happen to be the counterfactuals where the candidate lost in Florida. By contrast, across counterfactual worlds, there is a very low correlation between a Republican candidate winning in a Republican stronghold (like Texas) and the candidate winning the presidency – therefore we are unlikely to say that the candidate won the presidency *because* he won in Texas.

Second, the model provides a natural explanation to a well-established phenomenon: causal selection is influenced by the statistical *normality* of candidate causes. Researchers have found that the extent to which C is considered a cause of E depends on people’s perception of the prior probability of C, as well as the prior probability of other factors [1, 2, 5, 6, 7, 8]. Early theories held that, as a general rule, C is more likely to be selected as a cause when it is *abnormal* (i.e. rarer, or less expected) [6, 9]; according to this family of accounts, people tend to select the lit match over oxygen as the cause of a forest fire because a lit match is less normal than the presence of oxygen in the air. Yet subsequent research has found that the effect of normality on causal selection is more complex, because in some cases people are actually more likely to select normal over abnormal causes [1, 2, 7]. In a recent study which systematically investigated such effects, researchers have concluded that no current model of causal selection can fully account for the complex pattern of effects of statistical normality on causal selection [2]. Our model provides such an account. Furthermore, it outperforms other known computational

models of actual causal strength when tested against a high-resolution dataset of human graded causal judgments.

In a first section we present the intuition motivating the paper. Then we introduce a formal model of actual causal strength judgments. Third, we validate the model by assessing its fit to empirical data. In the discussion, we explore possible ways to extend the model, highlight aspects of actual causation judgments that the model is not designed to account for, and discuss the possible function of the human concept of actual causation.

1 Portable causal dependence

Consider someone who writes his birth date on a lottery ticket, and goes on to win the lottery. Under classical definitions of causality in philosophy and computer science (e.g., [10, 11]), the man won *because* he put his birth date on the ticket: had he not entered his birth date, he would not have won the lottery. Yet this causal assignment sounds odd: it seems to imply that entering one’s birth date is a good *general* strategy for winning the lottery. In other words, it seems like we read causal statements as containing general information, rather than being exclusively focused on the specific event under consideration.

Accordingly, researchers have argued that causal selection is designed to track causal dependence relations that are highly “portable”, meaning that they apply across a variety of situations [3, 4, 12]. Under this account, when we assign causes to an event, we favor causal relations that are not too dependent on the specific details of the current situation. For instance, we select the lit match as a cause of the forest fire, because across a variety of plausible background circumstances, striking a match inside a forest would result in a forest fire. But we do not explain the man’s lottery success by the fact that he

entered his birth date, since in most similar circumstances, entering his birth date would not have made him win the lottery.

Here we formalize this idea of “portable causal dependence” with a computational model of causal selection. We suggest that when someone assigns causes to an event E , she assigns a causal score to a candidate cause C by computing an estimate of the ‘average causal effect’ of C on E across situations somewhat similar to the one that she witnessed. She will then consider as genuine causes only candidate causes to which she assigned a sufficiently high causal score.

Below we informally explore how one can define the concept of an ‘average causal effect’ – we then formalize these ideas in the next section. For concreteness, we take the following situation as an example. One stands in a forest whose ground is covered in dry leaves, there is oxygen in the air, one strikes a match, and the forest catches fire. For each candidate cause (e.g. ‘the match is lit’, ‘the oxygen content of the air is 21%’, etc), we want to generate a score describing the ‘average causal effect’ of this variable on the intensity of the forest fire, across a variety of different situations.

Intuitively, an average causal effect is an average of *specific* causal effects. The specific causal effect of a candidate cause C on an outcome event E can be seen as the extent to which E causally depends on C in a particular situation. A popular way to define causal dependence, in cognitive science, philosophy and computer science, is in terms of counterfactuals: we imagine alternative worlds in which C has been set, by *fiat*, to a value different than its value in the current situation; E causally depends on C to the extent that, in such worlds, the value of E is different from its value in the actual world [1, 13, 14]. Therefore, a natural way to define the causal dependence of E on C in a specific situation is to assess what would happen to E , in this situation, if we changed

something about C , while keeping everything else about the situation constant, down to the tiniest details. For instance, in the situation described above, we can ask whether the intensity of the forest fire causally depends on the oxygen content of the air: if we changed the oxygen content of the air, by how much would the intensity of the fire change? The intensity of the forest fire causally depends on the oxygen content of the air to the extent that changes in oxygen content are associated with large changes in the intensity of the fire.

We can compute an average causal effect by computing the specific causal effect of C on E in many situations, including situations that differ from the situation one is making a causal judgment about. For instance, in addition to what actually happened, we also imagine situations where the ground is wet, or situations where no match has been lit; for each such situation, we compute the specific causal effect of C on E . It seems reasonable to imagine situations in proportion to how likely each situation is – for instance, we imagine situations where there is oxygen in the air with much higher frequency than situations where there is no oxygen in the air. The average causal effect of C on E is then simply the average of all specific causal effects of C on E across the situations we imagined.

Of course, different variables that have an influence on the occurrence of forest fires are measured in different units: oxygen content can be measured as the percentage of the air that is made of O_2 , while air temperature can be measured in Fahrenheit, or Celsius. If we are to compare their effects, we need to measure them in a common ‘unit-free’ currency. Empirical scientists face this problem all the time when they want to compare the strength of effects across different domains, and they typically solve it by converting variables into standard deviation units, for instance when they compute ‘effect size’ measures such as Cohen’s d or Pearson’s r . Similarly, we can standardize causal effect

measures by converting the relevant variables in standard deviation units. For instance, we can standardize our measure of oxygen content by dividing it by the standard deviation of oxygen content across all situations we imagined.

It is easy to see how such an algorithm for computing causal scores would assign a very low score to oxygen as an explanation of the forest fire. First, a forest fire is a relatively abnormal situation (most forests are *not* on fire) – this means that most of the situations we will imagine are situations without a forest fire (i.e. situations in which the intensity of the forest fire is 0). In such situations, changes to the oxygen content of the air will not on their own start a fire. Therefore, in most situations we imagine, the specific causal effect of oxygen is 0. Second, because the oxygen content of the air fluctuates only minimally across situations, only abnormally large deviations in the oxygen content of the air would produce meaningful changes in the intensity of a fire. In other words, a one-standard-deviation change in the oxygen content of the air will typically have a very small effect on the intensity of a fire; as a result, even in situations where the specific causal effect of oxygen is non-null, it will be small.

In summary, it seems that we can formalize the idea of portable causal dependence by formulating an algorithm which makes use of *nested counterfactuals*: counterfactuals within counterfactuals ([1] use a similar approach). We first imagine many counterfactuals to the situation we have witnessed; for each such counterfactual, we estimate the extent to which a change in the value of the candidate causal variable would result in a change in the outcome variable. Then we average all these estimates to generate a causal score. In the next section we formally define such an algorithm.

2 A formal model of actual causal strength

2.1 Structural equation models

In order to explain how people assign causes to events, we need to have a model of how the mind represents the causal structure of a situation. To that end, we use the formalism of structural equation models [13, 14]. In a structural equation model, the world is described in terms of variables. For example, the variable F can be used to represent whether the forest is on fire, where F takes the value 0 if the forest is not on fire, and 1 if it is (alternatively, one may treat F as a continuous variable representing the intensity of the fire). A variable is either endogenous or exogenous, depending on whether it is influenced by other variables in the model. The value of an endogenous variable is set deterministically as a function of other variables, as defined by a structural equation. Values of the exogenous variables are set randomly according to probability distributions specific to each such variable.

For example, we can model the causal system described in the forest fire scenario with one endogenous variable F , representing the state of the forest, and three exogenous variables: a variable L representing whether the match is lit, a variable D representing whether the ground is dry, and a variable Ox representing the amount of oxygen in the air. We assign a probability distribution to the value of each exogenous variable; here the average value of Ox would be high, with a relatively low standard deviation; there would be a low probability that $L = 1$, and a medium probability that $D = 1$. The state of the endogenous variable F is determined by the structural equation:

$$F = f(Ox, L, D)$$

where f is some function. In the simple case where we model all three exogenous variables as binary, the equation would be:

$$F = \min(Ox, L, D)$$

Such that F takes the value 1 if and only if there is oxygen in the air AND the match is lit AND the ground is dry.

In the structural equation formalism, an *intervention* is defined as the act of setting a variable V to a new value v of one’s choice. We do so by replacing the structural equation (or, in the case of an exogenous variable, the probability distribution) that would normally determine the variable’s value by a new equation $V = v$. An intervention can be seen as analogous to performing an experiment where we manipulate the value of the variable. Definitions of causality typically hold that a variable C has a causal influence on E if an intervention on C can have a downstream effect on the value of E [13, 14]. Intuitively, if changing the value of C does not have an effect on E , then C probably cannot count as a cause of E (and any correlation one may otherwise observe between the two variables must be due to a causal effect of E on C , or the causal influence of other variables).

Using the formalism of structural equation models, we can now formally state our model of actual causal strength.

2.2 Model

We refer to a specific state of a causal system as a ‘world’. To assess the causal effect of variable C on variable E , we use the following algorithm, operating

over a structural equation model of the causal system¹.

-Simulate a large number of worlds by sampling the set of possible worlds, according to the prior probabilities of the exogenous variables (in other words, sample worlds in proportion to how likely each world is). For each such world, the values of the endogenous variables are then determined naturally according to the structural equations. For each variable V in the causal system, compute the standard deviation σ_V of the variable value across all sampled worlds (for exogenous variables, this can simply be read off from the variable’s associated probability distribution).

-For each world generated that way, simulate a counterfactual ‘twin’ world by making an intervention on C , which sets C to a new, randomly sampled value. Then the values of the endogenous variables in this twin world are set naturally according to the structural equations.

-For each pair of worlds thus generated, compute the specific causal effect of C on E by taking the ratio of the change in the value of E to the change in the value of C between the two worlds ($\frac{\Delta E}{\Delta C}$), and multiplying this ratio by the standardizing factor $\frac{\sigma_C}{\sigma_E}$.

-The causal score of C on E is the average causal effect of C on E : the average of all specific causal effects across all pairs of worlds. Formally, we can denote it as $k_{C \rightarrow E}$ and write it as:

¹The process described here is best interpreted as belonging to Marr’s computational level of description [15]. That is, we are not suggesting that the human mind necessarily runs this exact algorithm (we ourselves did not run this sampling algorithm when implementing the model for assessing its fit to the human data – instead we relied on analytically derived exact solutions). Rather, we suggest that whichever algorithm the mind runs when computing a causal score, this algorithm’s output will approximate the output of the process described here.

$$k_{C \rightarrow E} = \frac{\sum_{i=1}^n (\frac{\Delta E}{\Delta C})_i \sigma_C}{n \sigma_E}$$

where n is the number of simulated world pairs.

The first step of the algorithm generates a large number of possible worlds, ensuring that we can look at the effect of C on E across a large number of different background circumstances, where these circumstances are represented in proportion to how likely they are to arise. The second step looks at each of these worlds in turn, asking about the strength of the causal dependence of E on C in each world. Our measure of causal dependence is standardized by the ratio of the standard deviation of C to the standard deviation of E . This standardization is akin to what scientists do when they compute statistical measures of effect size such as a Cohen’s d or a Pearson’s r ; it allows measures of causal effects to be unit-free, and thus comparable across domains. Formally, the standardization enables us to interpret a causal effect as the number of standard deviations by which the value of E changes as we make a one standard deviation change in the value of C . Finally, the last step of the algorithm takes the average of all the causal dependence scores computed in this way.

An elegant aspect of this model is that if C and E obey the “no-confounding assumption” [13], then $k_{C \rightarrow E}$ is simply the correlation between C and E across worlds sampled in step 1 (we prove this for the case of binary variables in the appendix). The “no-confounding assumption” holds when C has a causal influence on E , E does not have a causal influence on C , and no variable has a causal influence on both C and E . Intuitively, when this assumption holds, the relationship between C and E is not confounded by third variables, so we can read the causal effect of C on E from the correlation between C and E even in ‘observational’ data (i.e. data which was generated without performing any

intervention) [13].

3 Testing the empirical fit of the model

When judging whether a factor is causal, people seem to be sensitive to its statistical frequency, or its probability [1, 8, 16]. For instance, when a lit match and the presence of oxygen in the air conspire to start a forest fire, people judge that the match was the cause, and this might be because lit matches are more statistically abnormal than the presence of oxygen. Statistical normality is explicitly built into our model, since we assume that people sample counterfactual worlds in proportion to their respective prior probabilities. However, the form of the model was not explicitly motivated by the aim of explaining any specific effect of normality on causal judgment. Here we show that it nonetheless explains such effects.

We rely on a recent dataset collected by Morris et al. [2], which provides fine-grained data on human actual causality judgments. This study replicated several classic empirical findings about causal judgments, and uncovered several new patterns, thanks to its unprecedented high resolution.

Morris et al. asked participants to read the following vignette:

A person, Joe, is playing a casino game where he reaches his hand into two boxes and blindly draws a ball from each box. He wins a dollar if and only if he gets a green ball from the left box and a blue ball from the right box. Joe closes his eyes, reaches in, and chooses a green ball from the first box and a blue ball from the second box. So Joe wins a dollar.

Participants were asked the following question:

Please tell us how much you agree or disagree with this statement: Joe’s first choice (where he chose a green ball from the first box) caused him to win the dollar.

In a first experiment, participants saw the vignette shown above, which describes a conjunctive structure (Joe needs to draw a green ball from the first box, AND a blue ball from the second box, in order to win). In a second experiment, another set of participants read the same vignette, minimally modified so as to depict a disjunctive structure (Joe needs to draw a green ball from the first box, OR a blue ball from the second box, in order to win).

Participants were shown pictures of the two boxes. Across conditions, the experimenters systematically varied the proportion of green balls in the first box and blue balls in the second box. The proportion of green balls in the first box varied from 0.1 to 1, in 0.1 increments; the proportion of blue balls in the second box was similarly and independently manipulated. In total there were therefore 100 different possible states of the causal system. The dataset collected by Morris et al. therefore provides a rich, fine-grained empirical bedrock against which to test the quantitative predictions made by theories of human causal judgment. Morris et al. [2, 17] used their dataset to conduct such tests for the most prominent existing computational models of graded causal judgment [1, 9, 17, 18, 19, 20]. Their work thus makes two things possible. First it allows us to test the quantitative fit of our model. Second it allows us to compare this fit to the fit of other models. To anticipate, we find that our model fits the human data well, and overall does so significantly better than every model tested by Morris et al. [2, 17].

Following [17], we generate predictions for two versions of our model. The first version is the baseline version of the model. The second version is a “normalized” version, where we assume that when participants judge the causal importance of a draw of the green ball, their judgment incorporates an element of comparison to the causal importance of drawing the blue ball². This normalization is assumed to be accomplished by the softmax function:

$$\tilde{k}_{G \rightarrow D} = \frac{e^{k_{G \rightarrow D}}}{e^{k_{G \rightarrow D}} + e^{k_{B \rightarrow D}}}$$

Where $k_{G \rightarrow D}$ is the baseline causal strength ascribed to the draw of the green ball, and $k_{B \rightarrow D}$ is the baseline causal strength ascribed to the draw of the blue ball [17]. We also consider a baseline and a normalized version for all the models that are studied in Morris et al. [2, 17]. For each causal structure, we compute the predictions of our model by deriving analytical expressions corresponding to the correlation between “Joe draws a green ball” and “Joe wins a dollar” in the limit of an infinity of samples (see appendix for derivation). R code to reproduce analyses and figures is available in the electronic supplementary materials.

3.1 Results

3.1.1 Conjunctive structure

Human average ratings for the conjunctive structure are shown in figure 1, which plots participants’ average causal ratings for the green ball draw, for

²An important empirical motivation for normalization is the fact that intuitively, in a disjunctive causal structure, one would assign a greater causal score to a candidate cause when the other candidate cause is not present. See also [17] for additional arguments for normalization.

each combination of ball draw probabilities. Their ratings can range from 0 to 1, scaled from the original range of 1 to 9. The predictions of our model are shown in figure 2.

Figure 1: **Human judgments in the conjunctive structure.** The focal variable is ‘drawing green’ while the alternate variable is ‘drawing blue’.

Figure 2: **Judgments made by the normalized version of the model in the conjunctive structure.** The focal variable is ‘drawing green’ while the alternate variable is ‘drawing blue’.

First, we observe that both the human data and the model exhibit two well-known effects of statistical normality on causal judgment. The first effect is *abnormal inflation*: as “drawing green” (the focal variable) becomes less likely, causality ratings for “drawing green” increase [5, 6]. The second effect is *supersession* : as “drawing blue” (the alternate variable) becomes more likely, causality ratings for “drawing green” increase [8]. Second, both the data and the model show the same nonlinear effects: abnormal inflation is at its strongest when $Pr(\textit{green})$ goes from .9 to 1, and supersession is at its strongest when $Pr(\textit{blue})$ goes from .9 to 1.

Third, we studied the performance of each model by computing the item-level correlation between a model’s predictions and participants’ average causal ratings. The normalized version of our model achieved the best fit, $r(98) = .89$ ($p < .001$); the baseline version of our model, which achieved a correlation of $r(98) = .87$, had a marginally worse fit (Williams’ t-test, $t(97) = -1.88$, $p = .06$). The model by Morris et al. [17] made almost the same predictions as our model, and correlated with the human data at $r(98) = .87$ (normalized version, $p < .001$) and $r(98) = .86$ (baseline version, $p < .001$); nevertheless the normalized version of our model achieved a better fit than either version (Williams’ t-tests, $t(97) = 6.61$, $p < .001$ for the comparison with the normalized Morris model; $t(97) = 19.25$, $p < .001$ for the comparison with the baseline Morris model)³.

³The fact that a correlation of $r(98) = .89$ can be significantly higher than $r(98) = .87$ seems counter-intuitive. Such an outcome can occur because the predictions of the two models are extremely correlated with each other $r(98) = .997$, and this shared variance increases the power of Williams’ t-test. However, given the low magnitude of the difference, we hesitate to draw strong implications from its statistical significance. The fact that the two models make almost the same predictions can be explained analytically. In the appendix we show that, in this conjunctive causal structure, the causal score assigned by the Morris

Other models performed less well; out of them, the normalized versions of the Icard, Spellman and Delta-P / Power PC models (which make the exact same predictions as each other) performed the best, correlating with the human data at $r(98) = .79$, $p < .001$. Both versions of our model achieved a significantly better fit to the human data than these and all remaining models, all $ts > 3.43$ (all $ps < .001$).

3.1.2 Disjunctive structure

Human average ratings for the disjunctive structure are shown in figure 3, which plots participants' average causal ratings for the green ball draw, for each combination of ball draw probabilities. The predictions of our model are shown in figure 4.

Figure 3: **Human judgments in the disjunctive structure.** The focal variable is ‘drawing green’ while the alternate variable is ‘drawing blue’.

et al. model is equal to the square of the causal score assigned by our model (i.e. it is equal to the square of the correlation between “Joe draws a green ball” and “Joe wins a dollar”).

Figure 4: **Judgments made by the normalized version of the model in the disjunctive structure.** The focal variable is ‘drawing green’ while the alternate variable is ‘drawing blue’.

Both the human data and the model exhibit a well-known effect of statistical normality on causal judgment in disjunctive structures. This effect is *abnormal deflation* : as “drawing green” (the focal variable) becomes less likely, causality ratings for “drawing green” decrease [1, 7, 21]. They also exhibit a trend, *reverse supersession*, that had not been identified prior to the study by Morris et al. [2]: as “drawing blue” (the alternate variable) becomes less likely, causality ratings for “drawing green” increase. Both of these effects are the ‘flipped’ versions of the effects which occur in conjunctive causal structures.

Morris et al. [2] also highlight three interesting non-linear effects in the human data. First, when the alternate variable is certain ($Pr(blue) = 1$), the abnormal deflation effect disappears. Second, when the focal variable is certain ($Pr(green) = 1$), the reverse supersession effect disappears. Third, the reverse supersession effect is strongest as the alternate variable approaches certainty

($Pr(blue) \approx 1$). Our model partially reproduces these effects: the baseline version of the model exhibits the first and the third effect, but the normalized version of the model does not exhibit any of them (or does so weakly). In the appendix, we speculate that this may be because the exponentiation used in the softmax function is attenuating these non-linear effects; we report exploratory analyses showing that when using a simpler normalization function, which does not rely on exponentiation, our model reproduces the three non-linear effects found in the human data. This suggests that people may be using a slightly different normalization function than the softmax function we use here.

We studied the quantitative fit of each model by computing the item-level correlation between a model’s predictions and participants’ average causal ratings. The best performing models were the normalized version of our model ($r(98) = .64, p < .001$), the Icard model (baseline version, $r(98) = .64, p < .001$; normalized version, $r(98) = .62, p < .001$), and the normalized Delta-P model ($r(98) = .62, p < .001$). None of these four models fit the data better than any other, all $ts < .73$, all $ps > .47$. The next best model was the baseline version of our model ($r(98) = .47, p < .001$), which performed less well than the models above (all $|ts| > 3.69$, all $ps < .001$), but better than all other models (all $ts > 4.03$, all $ps < .001$).

3.2 Discussion

In summary, the normalized version of our model fits the data better than all other models in the conjunctive structure, although it is almost tied with the model by Morris et al. [17]. In the disjunctive structure, it is one of a few models tied for first place, along with the Delta-P and the Icard model.

Because it fits the data better than the Icard model and the Delta-P model in the conjunctive structure scenario, and better than the Morris model in the disjunctive scenario⁴, it is the model which is most consistently able to account for human behavior across the two tasks.

From a qualitative point of view, it is also the model which best reproduces the general pattern found in the human data, sharing the same four general trends (supersession and abnormal inflation in conjunctive structures, reverse supersession and abnormal deflation in disjunctive structures) along with more subtle non-linear effects. There is however one general feature of the human data which neither our model, nor any of the other models tested, emulate: people tend to give much higher causal ratings in the disjunctive structure ($M = .70$) than the conjunctive structure ($M = .47$), $p < .001$. It also remains an open question why our model has a better fit in the conjunctive ($r = .89$) than in the disjunctive structure ($r = .64$).

4 General discussion

We have proposed that as part of the process of judging the extent to which C caused E, people compute a measure of the average causal effect of C on E across counterfactual states of the causal system. This model formalizes the proposal that people preferentially select “portable” causal explanations [3, 4],

⁴The model by Morris et al. [17] predicts causality ratings to be uniformly 0 in the disjunctive causal structure, because it requires causes to be necessary. In the disjunctive causal structure scenario, Joe draws both a green ball and a blue ball, so neither the green ball nor the blue ball was necessary for Joe to win the game. For completeness, we also considered a version of their model which relaxes the requirement that causes be necessary; its predictions were negatively correlated with human ratings, $r(98) = -.26$ for the baseline version, $r(98) = -.63$ for the normalized version.

and can account for the complex pattern of effects of statistical normality on human graded causal judgments.

Now we highlight opportunities to extend the model, and discuss its intended scope. First, when sampling counterfactuals, people probably do not take only *statistical* normality into account – their notion of normality is influenced by other factors. Second, our measure of causal effect mostly captures the monotonic causal effects that a variable has on another variable, and might miss some more subtle effects (for instance, U-shaped relationships between two variables). Third, our model is not meant as a complete theory of actual causation judgment; intuitively, whether C is judged as having caused E should not depend only on whether C is highly causal for E in general; it should also depend on what happened in the actual world. We discuss these points in turn; we end by addressing the possible function of actual causation judgments.

4.1 Non-statistical normality

The statistical normality of a candidate cause (its frequency, or its probability) plays an important role in our model, because the sampling algorithm we propose samples counterfactual worlds in proportion to how likely they are. However, normality does not always have a statistical meaning [22]. Normality can be prescriptive: violating a rule can be considered less normal than upholding that rule, even when most people do the former. Normality can also have a functional meaning: for instance an engine that fails behaves abnormally with respect to the designer’s intentions about its function. All three kinds of normality have effects on people’s causal judgments [12], and these effects seem to be similarly patterned: for instance, for both statistical and prescrip-

tive normality, abnormal inflation occurs in conjunctive causal structures, and abnormal deflation occurs in disjunctive structures [1].

Our model focuses on statistical normality, because it is the easiest factor to manipulate precisely enough to allow fine-grained tests of computational models (as in [2]). One could extend our model by following Icard et al. [1], and assume that people sample counterfactual worlds in proportion to how normal each counterfactual world is, where normality is determined by a mix of statistical, prescriptive and functional considerations. Of course, a major challenge would then be to specify exactly how these three different criteria for normality are combined to yield a single normality rating (but see [23] for hints of a possible solution).

4.2 Non-monotonic causal effects

Our measure of average causal effect relies on a simple average of the causal effects computed for each counterfactual world. This simple procedure may paint a misleading picture in causal systems where the influence of a variable on another variable is non-monotonic. For instance, if the influence of C on E follows an inverted-U-shape (such that increases in C cause increases in E at low levels of C, but decreases in E at high levels of C), then averaging the positive and the negative effects of C might make them cancel each other, resulting in a measure of average causal effect close to 0. Future work could attempt to develop measures that are more sensitive to such patterns.

4.3 The importance of what actually happened

Our model sheds light on an important aspect of actual causation judgment, but is not meant as a complete theory of actual causation. To see why, con-

sider for instance the following scenario.

Jake grows plants in his greenhouse. A plant's adult height is influenced in equal degrees by the amount of water and the amount of sunlight it receives. Indeed, from careful observations, Jake has concluded that a one standard deviation change in the amount of water received has the same average effect on the plant's growth as a one standard deviation change in sunlight exposure. There is a plant in Jake's greenhouse which has received a lot of water, but barely any sunlight. It is very tall.

Water and sunlight have the same average causal effect in this scenario, yet we clearly have the intuition that this specific plant is tall because of the water it received, not because of the sunlight it was exposed to. This suggests that the exact value of a causal variable *in the actual world* matters to actual causation judgment.

Consider also the following scenario (inspired by [17]).

A college's main admission criterion is the SAT scores of an applicant. However, Julia has a stellar athletic record, so the admission office does not even bother looking at her (excellent) SAT scores and she gets accepted.

Clearly, Julia's SAT scores are not a cause of her admission, even though, in general, SATs are the best predictor of a candidate's admission at this college. It seems that for a candidate cause to count as a cause, there has to be a causal pathway via which, *in the actual world*, it influences the outcome variable.

Such cases, along with the empirical success of our model involving only general normality considerations, suggest that people’s judgments about what caused a specific event may be the output of a process involving two complementary sets of algorithms. While the first set of algorithms is very sensitive to what actually happened, the second set of algorithms is mostly sensitive to the average causal effect of a variable across counterfactual worlds. Our model falls into the second category of algorithms.

A good reason to posit the existence of two different sets of algorithms is that existing models of actual causation judgment seem to fall into two distinct types. Some theories model actual causation as being mostly shaped by general normality considerations [9, 1, 17]; other theories ignore normality considerations and are sensitive to which events did and did not occur in the actual world [24, 25, 26, 11, 27]. Each class of models has its own strengths and weaknesses. Models centered around general normality considerations typically excel at capturing patterns of graded judgments in simple tasks; however, they fail at reproducing human intuitions in many thought experiments.⁵ The other class of models handle such thought experiments well, but they typically deliver only simple Yes/No judgments, so they cannot account for the richness of graded judgments made by humans. This complementarity suggests that both kinds of models are necessary to provide a full picture of how human actual causation judgment works.

⁵For instance, Icard et al. [1] acknowledge that their model gives good predictions of human intuitions in simple causal structures, but that it would need additional assumptions in order to accommodate more complex scenarios. Their model, just like the present model, assigns a high causal rating to Julia’s SAT scores in the college admission scenario described above. The model presented in [17], though it provides the correct intuition to the college admission problem, fails to account for ‘overdetermination’ cases, where two causes jointly cause an outcome, but either one in isolation would have been sufficient.

4.4 The function of actual causation

The present model is a formalization of a verbal idea about the *function* of the human concept of actual causation – namely the idea that causal ascriptions are designed to identify explanations that would generalize across a wide range of background circumstances [3, 4]. The tight fit between the model’s predictions and human causal intuitions suggests that a functional approach to actual causation is a promising one.

Yet from a functional perspective, the concept of actual causation is puzzling. Modelers typically assume that people derive causal judgments from a pre-existing internal representation of the causal structure of the world; for instance, our model assumes that the agent making causal judgments possesses a structural equation model representing the situation. The structural equation model seems to already contain all the relevant information: it specifies the value of each variable in the actual world, the causal relationships between them, as well as the probability distributions governing the values that variables would take across different contexts. An agent possessing such an internal causal map of the situation would be perfectly able to perform useful operations (answer counterfactual questions, generate predictions, etc), *even if he did not have a concept of actual causation*. In other words, it seems that an algorithm for generating a causal score would not provide any new information, beyond that which is already contained in the structural equation model, that would help the agent make better decisions. Why bother assigning causes to events, if this cannot help us think better?

One possibility is that the function of actual causation is not to help us *think* better, but to help us *communicate* better⁶. If we want to explain the

⁶A broader version of the present hypothesis is that the generation of a causal score is an

origins of the forest fire to a friend, we could try to enunciate all the details of our internal causal map of the situation, but this would take too long (e.g., “There was oxygen in the air; if there hadn’t been oxygen in the air, the fire would not have started. The ground was dry; if the ground had been wet...”). Instead, saying that the fire was caused by a lit match allows us to efficiently convey the most essential information that is contained in our internal causal map. If this hypothesis is correct, a functional analysis of the kind of causal information that people would most need to transmit to others may be a rich source of new hypotheses about the grammar of actual causation.

Our findings are also consistent with a recent proposal that causal selection is designed to select good interventions [17]. According to this proposal, causal selection is designed to identify our best-bet intervention if we want to bring about an outcome but do not know the exact state of the causal system. Under the model proposed in the present paper, the causal score of C on E, $k_{C \rightarrow E}$, provides a good measure of intervention effectiveness. It tells us that, on average, we can expect that intervening on C will result in a change of $k_{C \rightarrow E}$ standard deviation units in E for each one standard deviation unit change in C. Therefore, if we do not know the current state of the causal system, but want to set E to a certain value, we are generally better off making an intervention on the variable X with the highest $k_{X \rightarrow E}$. For instance, in the simple conjunctive structure used in the experiments described above, intervening to make Joe instance of ‘representational exchange’ [28]: it takes information encoded in a given format (here, information encoded in a structural equation model) and converts it to a format that can be widely shared with cognitive systems that cannot access the original content. These target cognitive systems may be situated in other minds, but can also reside in the same mind that hosts the original content. In other words, causal scores may be summaries aimed at cognitive systems that cannot read the causal map directly (see [29, 30, 31, 28] for general considerations about information transfer between different parts of the mind).

draw the ball with the highest causal score is our best bet if our goal is to make Joe win a dollar. In the disjunctive structure, intervening to prevent Joe from drawing the ball with the highest causal score is our best bet if we want to prevent Joe from winning a dollar.⁷

5 Conclusion

People attribute causes to events so spontaneously and effortlessly that it is easy to forget that when they do so, they solve the daunting problem of selecting a few causes out of the massive set of all factors contributing to that event. Here we have proposed that, as part of this process, they attribute a causal score to a candidate cause by computing a measure of its average causal effect across counterfactual worlds. If a simple assumption (the ‘no-confounding assumption’) holds, then this can be achieved by simply computing the correlation, across counterfactual worlds, between the candidate cause and the outcome. This simple idea can account for a wide range of human intuitions about causality.

⁷If our goal is to prevent Joe from winning in the conjunctive structure, or make him win in the disjunctive structure, then an intervention on any ball will do.

References

- [1] Icard TF, Kominsky JF, Knobe J. Normality and actual causal strength. *Cognition*. 2017;161:80–93.
- [2] Morris A, Phillips J, Gerstenberg T, Cushman F. Quantitative causal selection patterns in token causation. *PloS one*. 2019;14(8):e0219704.
- [3] Lombrozo T. Causal–explanatory pluralism: How intentions, functions, and mechanisms influence causal ascriptions. *Cognitive Psychology*. 2010;61(4):303–332.
- [4] Hitchcock C. Portable causal dependence: A tale of consilience. *Philosophy of Science*. 2012;79(5):942–951.
- [5] Hilton DJ, Slugoski BR. Knowledge-based causal attribution: The abnormal conditions focus model. *Psychological review*. 1986;93(1):75.
- [6] Kahneman D, Miller DT. Norm theory: Comparing reality to its alternatives. *Psychological review*. 1986;93(2):136.
- [7] Gerstenberg T, Icard T. Expectations affect physical causation judgments. *Journal of Experimental Psychology: General*. 2019;.
- [8] Kominsky JF, Phillips J, Gerstenberg T, Lagnado D, Knobe J. Causal superseding. *Cognition*. 2015;137:196–209.
- [9] Halpern JY, Hitchcock C. Graded causation and defaults. *The British Journal for the Philosophy of Science*. 2015;66(2):413–457.
- [10] Lewis D. Causation. *The journal of philosophy*. 1973;70(17):556–567.

- [11] Halpern JY, Pearl J. Causes and explanations: A structural-model approach. Part I: Causes. *The British journal for the philosophy of science*. 2005;56(4):843–887.
- [12] Hitchcock C, Knobe J. Cause and norm. *The Journal of Philosophy*. 2009;106(11):587–612.
- [13] Pearl J. Causality: models, reasoning and inference. vol. 29. Springer; 2000.
- [14] Halpern JY. Actual causality. MIT Press; 2016.
- [15] Marr D. Vision: A computational investigation into the human representation and processing of visual information. vol. 2. Inc., New York, NY; 1982.
- [16] Harinen T. Normal causes for normal effects: Reinvigorating the correspondence hypothesis about judgments of actual causation. *Erkenntnis*. 2017;82(6):1299–1320.
- [17] Morris A, Phillips J, Icard T, Knobe J, Gerstenberg T, Cushman F. Judgments of actual causation approximate the effectiveness of interventions. *Psyarxiv*. 2018;.
- [18] Cheng PW. From covariation to causation: a causal power theory. *Psychological review*. 1997;104(2):367.
- [19] Jenkins HM, Ward WC. Judgment of contingency between responses and outcomes. *Psychological monographs: General and applied*. 1965;79(1):1.
- [20] Spellman BA. Crediting causality. *Journal of Experimental Psychology: General*. 1997;126(4):323.

- [21] Henne P, Niemi L, Pinillos Á, De Brigard F, Knobe J. A counterfactual explanation for the action effect in causal judgment. *Cognition*. 2019;190:157–164.
- [22] Bear A, Knobe J. Normality: Part descriptive, part prescriptive. *Cognition*. 2017;167:25–37.
- [23] Bear A, Bensinger S, Jara-Ettinger J, Knobe J, Cushman F. What comes to mind? *Cognition*. 2020;194:104057.
- [24] Hitchcock C. The intransitivity of causation revealed in equations and graphs. *The Journal of Philosophy*. 2001;98(6):273–299.
- [25] Woodward J. *Making things happen: A theory of causal explanation*. Oxford university press; 2003.
- [26] Weslake B. A partial theory of actual causation. *British journal for the philosophy of science*. 2015;.
- [27] Halpern J. A modification of the Halpern-Pearl definition of causality. In: *Twenty-Fourth International Joint Conference on Artificial Intelligence*; 2015. .
- [28] Cushman F. Rationalization is rational. *Behavioral and Brain Sciences*. 2018;p. 1–69.
- [29] Fodor JA. *The modularity of mind*. MIT press; 1983.
- [30] Dehaene S, Naccache L. Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Cognition*. 2001;79(1-2):1–37.

[31] Barrett HC. Enzymatic computation and cognitive modularity. *Mind & language*. 2005;20(3):259–287.

6 Appendix

6.1 Proof that the causal metric reduces to a correlation coefficient under the no-confounding assumption

Let X and Y be two binary variables (with possible values 0 and 1), which obey the no-confounding assumption; that is:

$$Pr(y|do(x)) = Pr(y|x)$$

for all possible values x and y of X and Y , respectively [13]. $Pr(y|do(x))$ denotes the probability of $Y = y$ given that we have performed an intervention setting X to x [13].

We show that the causal metric $k_{X \rightarrow Y}$ is equivalent to the correlation between X and Y across randomly drawn, independent counterfactual worlds.

We have:

$$k_{X \rightarrow Y} = b_{Y,X}^k \frac{\sigma_X}{\sigma_Y}$$

where $b_{Y,X}^k = \frac{\sum_{i=1}^{n^*} (\frac{\Delta Y}{\Delta X})_i}{n^*}$, and n^* is the number of world pairs where $\Delta X \neq 0$.

With binary variables, we can classify the world pairs generated by the sampling-resampling algorithm according to the three following categories:

- world pairs where $\Delta X = 0$,
- world pairs where $\Delta X = 1$,
- world pairs where $\Delta X = -1$

The first kind of world pair is ignored for the computation of $b_{Y,X}^k$, since $\frac{\Delta Y}{\Delta X}$ is undefined when $\Delta X = 0$. We denote the number of world pairs of the second kind as l , and the number of world pairs of the third kind as m .

World pairs of the second kind are made of worlds where $X = 0$ in the original world and $X = 1$ in the twin world. Therefore the proportion of twin worlds with $Y = 1$ is equal to $P(Y = 1|do(X = 1))$, while the proportion of original worlds with $Y = 1$ is equal to $P(Y = 1|X = 0)$. It is easy to see that $\frac{\sum_{i=1}^l (\Delta Y)_i}{l}$ must be equal to the difference between the proportion of twin worlds with $Y = 1$ and the proportion of original worlds with $Y = 1$. Therefore we have

$$\frac{\sum_{i=1}^l (\Delta Y)_i}{l} = Pr(Y = 1|do(X = 1)) - Pr(Y = 1|X = 0)$$

i.e.

$$\frac{\sum_{i=1}^l (\frac{\Delta Y}{\Delta X})_i}{l} = Pr(Y = 1|do(X = 1)) - Pr(Y = 1|X = 0)$$

World pairs of the third kind are made of worlds where $X = 1$ in the original world and $X = 0$ in the twin world. Therefore the proportion of twin worlds with $Y = 1$ is equal to $Pr(Y = 1|do(X = 0))$, while the proportion of original worlds with $Y = 1$ is equal to $Pr(Y = 1|X = 1)$. $\frac{\sum_{i=1}^m (\Delta Y)_i}{m}$ is equal to the difference between the proportion of twin worlds with $Y = 1$ and the proportion of original worlds with $Y = 1$. Therefore we have

$$\frac{\sum_{i=1}^m (\Delta Y)_i}{m} = Pr(Y = 1|do(X = 0)) - Pr(Y = 1|X = 1)$$

i.e.

$$\begin{aligned} \frac{\sum_{i=1}^m (\frac{\Delta Y}{\Delta X})_i}{l} &= -(Pr(Y = 1|do(X = 0)) - Pr(Y = 1|X = 1)) \\ &= Pr(Y = 1|X = 1) - Pr(Y = 1|do(X = 0)) \end{aligned}$$

By the no-confounding assumption, we have

$$Pr(Y = 1|do(X = 1)) - Pr(Y = 1|X = 0) = Pr(Y = 1|X = 1) - Pr(Y = 1|do(X = 0))$$

$$= Pr(Y = 1|X = 1) - Pr(Y = 1|X = 0)$$

i.e. the average $\frac{\Delta Y}{\Delta X}$ is the same in world pairs of the second and third kinds.

$b_{Y,X}^k$ is the average $\frac{\Delta Y}{\Delta X}$ in the world pairs of the second and third kind:

$$b_{Y,X}^k = Pr(Y = 1|X = 1) - Pr(Y = 1|X = 0)$$

This is simply the regression coefficient $b_{Y,X}$ of Y on X ; that is, we have $b_{Y,X}^k = b_{Y,X}$. The correlation coefficient $r_{Y,X}$ between two variables X and Y is related to the regression coefficient $b_{Y,X}$ of Y on X via the following formula:

$$r_{Y,X} = b_{Y,X} \frac{\sigma_X}{\sigma_Y}$$

Therefore we have:

$$r_{Y,X} = b_{Y,X}^k \frac{\sigma_X}{\sigma_Y} = k_{X \rightarrow Y}$$

6.2 Analytical expressions for actual causal strength in simple conjunctive and disjunctive causal structures

Here we derive the analytical expressions for our actual causal strength score, in the simple causal structures used in Morris et al. [2]. The variable G takes the value 1 if Joe draws a green ball from the left box, 0 if he draws a non-green ball from the left box. Similarly, B is the variable representing whether Joe draws a blue ball from the right box, and D represents whether Joe wins a dollar. In both causal structures, the relation between “Joe draws a green ball from the left box” and “Joe wins a dollar” obeys the no-confounding assumption, since there is no variable with a causal influence on them both. Therefore, the causal score $k_{G \rightarrow D}$ simply corresponds to the correlation between “Joe draws a green ball from the left box” and “Joe wins a dollar”

across sampled counterfactual worlds. Here we compute these correlations in the limit of an infinity of samples.

The correlation coefficient $r_{Y,X}$ between two variables X and Y is related to the regression coefficient $b_{Y,X}$ of Y on X via the following formula:

$$r_{Y,X} = b_{Y,X} \frac{\sigma_X}{\sigma_Y}$$

Therefore we have:

$$r_{D,G} = b_{D,G} \frac{\sigma_G}{\sigma_D}$$

The variance of a random binary variable X is $Pr(X = 1)(1 - Pr(X = 1))$. For conciseness, we will denote probabilities using the shorthands $g = Pr(G = 1)$; $b = Pr(B = 1)$; and $d = Pr(D = 1)$. Therefore we have:

$$\sigma_G = \sqrt{g(1 - g)}$$

and

$$\sigma_D = \sqrt{d(1 - d)}$$

This yields:

$$r_{D,G} = b_{D,G} \sqrt{\frac{g(1 - g)}{d(1 - d)}}$$

We will use this along with the fact that, in the case of binary variables, we have:

$$b_{D,G} = Pr(D = 1|G = 1) - Pr(D = 1|G = 0)$$

(intuitively, the regression coefficient quantifies how much the probability that $D = 1$ changes for each unit change in the value of G).

6.2.1 Conjunctive causal structure

In the conjunctive causal structure, drawing a green ball from the left box instead of not drawing a green ball (or vice-versa) makes a difference to the

outcome only if one has drawn a blue ball from the right box. Therefore we have:

$$b_{D,G} = Pr(D = 1|G = 1) - Pr(D = 1|G = 0) = Pr(B = 1) = b$$

We also have

$$\begin{aligned} d &= Pr(G = 1 \wedge B = 1) \\ &= gb \end{aligned}$$

Combining these, we have:

$$r_{D,G} = b \sqrt{\frac{g(1-g)}{gb(1-gb)}}$$

i.e.

$$\begin{aligned} r_{D,G} &= \sqrt{\frac{gb^2(1-g)}{gb(1-gb)}} \\ r_{D,G} &= \sqrt{\frac{(1-g)b}{(1-gb)}} \end{aligned}$$

Note that the causal score derived in [17] for this causal structure is:

$$\frac{(1-g)b}{(1-gb)}$$

which is simply $r_{D,G}^2$.

6.2.2 Disjunctive causal structure

In the disjunctive causal structure, drawing a green ball from the left box instead of not drawing a green ball (or vice-versa) makes a difference to the outcome only if one has *not* drawn a blue ball from the right box. Therefore we have:

$$b_{D,G} = Pr(D = 1|G = 1) - Pr(D = 1|G = 0) = 1 - Pr(B = 1) = 1 - b$$

We also have

$$\begin{aligned} d &= Pr(G = 1 \vee B = 1) \\ &= g + b - gb \end{aligned}$$

Combining these, we have:

$$\begin{aligned} r_{D,G} &= (1 - b) \sqrt{\frac{g(1 - g)}{d(1 - d)}} \\ &= (1 - b) \sqrt{\frac{g(1 - g)}{(g + b - gb)(1 - g - b + gb)}} \\ &= \sqrt{\frac{(1 - b)g(1 - b)(1 - g)}{(g + b - gb)(1 - g - b + gb)}} \\ &= \sqrt{\frac{(1 - b)g(1 - g - b + bg)}{(g + b - gb)(1 - g - b + gb)}} \end{aligned}$$

i.e.

$$r_{D,G} = \sqrt{\frac{(1 - b)g}{g + b - gb}}$$

6.3 Non-linear patterns in the disjunctive causal structure.

Morris et al [2] report the three following non-linear effects in the human data for the disjunctive causal structure (see figs. 5 to 7):

- a)** When the alternate variable is certain ($Pr(blue) = 1$), abnormal deflation ceases to occur
- b)** When the focal variable is certain ($Pr(green) = 1$), reverse supersession ceases to occur
- c)** Reverse supersession is at its most powerful as the alternate variable approaches certainty (as $Pr(blue)$ approaches 1).

Figure 5: **Effect a.** Abnormal deflation ceases to occur when the alternate is certain.

Figure 6: **Effect b.** Reverse supersession ceases to occur when the focal is certain.

Figure 7: **Effect c.** Reverse supersession occurs mostly as the alternate variable approaches certainty (as $Pr(blue)$ approaches 1).

In what follows, we give an informal discussion of how these effects may be explained by our model, and why they suggest that people use a slightly different normalization procedure than the one we test in the main text.

Effects a) and c) follow naturally from the correlation model:

Explanation of (a): When $Pr(blue) = 1$, then Joe always wins a dollar no matter whether he draws a green ball or not. Therefore, the frequency of green balls in the box does not matter to his likelihood of winning: the correlation between ‘Joe draws a green ball from the left box’ and ‘Joe wins a dollar’ across counterfactual worlds is always 0, no matter the value of $Pr(green)$ (except when $Pr(green) = 1$). As a result, abnormal deflation disappears when $Pr(blue) = 1$.

Explanation of (c): When $Pr(blue)$ approaches 1, the causal effect of green drops to 0 (for the reason just outlined above), no matter the value of $Pr(green)$. Therefore, reverse supersession is at its most powerful when $Pr(blue)$ approaches 1.

Effect (b) makes sense if we assume that people normalize the causal effect of green by comparing it to the causal effect of blue. When $Pr(\text{green}) = 1$, the correlation between “Joe draws a blue ball from the right box” and “Joe wins a dollar” is 0, so the causal effect of blue is 0, no matter the value of $Pr(\text{blue})$ (except when $Pr(\text{blue}) = 1$). When the causal effect of blue is 0, then the normalization procedure always attribute 100% of the overall causal effect to green. As a result, reverse supersession disappears when $Pr(\text{green}) = 1$.

In other words, the three effects can be explained by the fact that when one variable is certain, the correlation between the other variable and the outcome is 0, so the unnormalized causal effect of that variable is 0. Indeed, the unnormalized version of our model exhibits effects (a) and (c) (see figs. 8 and 9). However, note that these explanations only work if we assume that the normalization procedure treats a causal effect of 0 as a ‘true zero’, and does not convert it into another number before comparing it to the causal effect of other variables.

Figure 8: **Effect a in the predictions of the baseline model.** Abnormal deflation ceases to occur when the alternate is certain.

Figure 9: **Effect c in the predictions of the baseline model.** Reverse supersession occurs mostly as the alternate variable approaches certainty.

Yet the normalization we use (following [17]) relies on a softmax function, which exponentiates the raw causal effect of each variable before comparing them to each other. By doing so it transforms 0s into 1s (since $e^0 = 1$). As a result, when we examine the normalized version of our model, the three non-linear effects we find in the human data are absent, or weak (see figs. 10 to 12).

Figure 10: (absence of) Effect a in the normalized model.

Figure 11: (absence of) Effect b in the normalized model.

Figure 12: **Effect c in the normalized model.**

By contrast, consider the following normalization function, which does not rely on exponentiation, and therefore treats zeroes as ‘true zeroes’:

$$\tilde{k}'_{G \rightarrow D} = \frac{k_{G \rightarrow D}}{k_{G \rightarrow D} + k_{B \rightarrow D}}$$

When using this normalization function, our model exhibits all three non-linear effects described above (see figs. 13 to 15).

Figure 13: **Effect a in the model with alternative normalization.** Abnormal deflation ceases to occur when the alternate is certain.

Figure 14: **Effect b in the model with alternative normalization.** Reverse supersession ceases to occur when the focal is certain.

Figure 15: **Effect c in the model with alternative normalization.** Reverse supersession occurs mostly as the alternate variable approaches certainty (as $Pr(blue)$ approaches 1).

In the empirical data, people very rarely ascribe causal scores of 0, so they are probably not using the exact normalization function shown above. One may speculate that they use a normalization procedure which treats raw causal

scores of zero as ‘true zeroes’ before comparing them to other raw causal scores, but then converts them into another number at the end of the procedure.