Causal judgments approximate the effectiveness of future interventions

Adam Morris¹, Jonathan Phillips¹, Thomas Icard², Joshua Knobe³,
Tobias Gerstenberg² & Fiery Cushman¹

¹Harvard University, ²Stanford University, ³Yale University

1 Abstract

When many things contributed to an outcome, people consistently judge certain ones to be more causal than others. For instance, people believe that a fire was more caused by the lit match than by the surrounding oxygen that fueled it. Why? Here, we offer a functional account of such patterns in causal judgment: By selecting causes as people naturally do, repeated judgments of whether something (e.g. the match) was the cause of an outcome (e.g. the fire) can be averaged to obtain the probability that introducing those things would produce the outcome (e.g., that lighting a match would start a fire). In other words, token causal judgments accumulate evidence about the general effectiveness of potential future interventions. We offer a formal account of this process, and show how it explains three basic qualitative features of causal judgment: why the causes people select tend (1) to be necessary, (2) to be abnormal, and (3) to lack abnormal counterparts.

Keywords: cause; causal selection; planning; intervention; sampling approximation

2 Introduction

Imagine that an ancient forest burns down. What caused the fire? In certain simple cases, we would all agree on the answer. If a careless camper dropped a lit match on dry leaves, we would judge her the cause. If the leaves were struck by lightning, we would point to the lightning strike.

But, on reflection, a multitude of factors contributed to the fire. The lack of rain; the National Park Service’s decision to allow campers in this forest; the store owner selling the match to the camper; the oxygen in the air—all of these were critical antecedents to the blaze, and yet people would not typically consider them to be the actual causes of the event. (Imagine a news anchor proclaiming that a forest fire was caused by oxygen in the air; it seems absurd.) Dry weather and oxygen recede in our minds, while campers and lightning stand out as the real causes of the fire.

This phenomenon is widespread: When multiple variables contributed to an outcome, humans judge certain variables more of a cause than others (Hart & Honoré, 1959/1985; Hitchcock & Knobe, 2009; Mackie, 1974). This process of identifying certain events as the true, relevant causes—sometimes called “causal selection” (Hesslow, 1988), or “specific” causation (Sloman, Barbey, & Hotaling, 2009; Walsh & Sloman, 2011)—pervades human thought and discourse.

Its psychological function, however, remains mysterious. When making token causation judgments (i.e., causal judgments of specific “token” events¹), why do people consistently foreground certain variables as more causal than others? (Cheng & Novick, 1991; Hesslow, 1988; Hilton, 1990)? Why not treat all contributing variables equally?

In this paper, we offer a functional account of certain key patterns in token causal judgments. We propose that token causal judgments—including people’s tendency to judge some variables more causal than others—are designed in part to assist in future

---

¹“Token causation” here means judgments of singular instances—e.g. judging whether Joe’s smoking habit caused his lung cancer (sometimes also called “singular” or “actual” causation see, e.g. Danks, 2017; Halpern & Pearl, 2005). This is opposed to “type causation”—e.g. judging whether smoking causes lung cancer in general. In this paper, we focus only on token causation judgments. Also, we sometimes describe variables (e.g. `lit match` and `oxygen`) as causes, but other times we describe events as causes. We treat these as equivalent; for our purposes, the variable `lit match` is a cause just in case the event of someone lighting a match is a cause.

decision-making by identifying effective ways to produce an outcome. Concretely, judging the extent to which a lit match was the cause of a fire in past situations helps approximate the probability that lighting a match, in future situations, would produce a fire. We derive an algorithm that accomplishes this goal, accumulating evidence from past causal judgments to guide future planning. Then, we show how some basic features of this algorithm capture important qualitative properties of causal judgment that are already well-documented, and provide a uniquely close fit to recent quantitative data. Taken together, our findings suggest that the token causal judgment is designed, in part, to accumulate evidence across diverse cases in order to approximate the general effectiveness of future interventions.

Our proposal is organized around two key claims: (1) token causal judgment is designed to estimate the effectiveness of future interventions, and (2) it accomplishes this in a cognitively efficient manner by employing a “sampling approximation”.

Causal judgments identify effective interventions Our model is inspired by recent philosophical treatments of token causation. They propose that people select causes that are “effective interventions”—i.e., good actions to perform if you want to actually bring about an effect (Hitchcock, 2012; Hitchcock & Knobe, 2009; Woodward, 2003)² For instance, they propose that there is an important connection between how a golfer judges the *cause* of a hole-in-one (“I put just the right spin on it!”) and how a golfer tries to *make* a hole-in-one (“I’d better put just the right spin on it!”). Judging what caused an outcome, and intervening to produce the outcome, are fundamentally connected.

We build upon this idea by arguing for a specific ordering of these judgments—a flow of information from one to the other. In principle, we might select token causes based on our prior understanding of effective interventions. In other words, a golfer might determine that the spin “caused” her hole-in-one because she knows that spin is, in general, a good way to make a hole-in-one. But we propose precisely the opposite

---

²The notion of “effective intervention”, as it is used here, is not the purely formal notion of intervening on a causal model in order to infer, define or reason about causal relations (Pearl, 2000). Rather, it is a more practical notion of intervention as an action performed to bring about an effect in the real world; a notion used to plan and make decisions based on a causal model that one has already defined.

relationship—that causal judgments accumulate knowledge about intervention effectiveness. In other words, we know how to make a good shot because of our past history of judging the causes of good shots.

There is something very natural about this idea: That, when we need to make something happen in the present, we ask, “What has tended to cause that thing in the past?” Yet, this kind of extrapolation is not trivial. Our primary contribution is to offer a formal model of how it can work—the algorithm we use to identify variables like “lots of spin” as the cause of a specific shot’s success, and the reason that this algorithm helps us identify successful interventions in the future. Notably, the algorithm that accomplishes this goal naturally leads to three very salient features of human causal judgment that have long been documented, but have previously lacked a normative justification. It also provides a uniquely good fit to recent behavioral data that quantifies the precise form of these qualitative features.

Causal judgments as an efficient sampling approximation Of all the ways of identifying optimal interventions, why employ causal judgments of past events? Even if such an approach could work, it is not obvious why it would be favored.

We propose that this method is favored because of its cognitive efficiency. In complex real-world settings, identifying effective interventions is difficult (cf. Eiter & Lukasiewicz, 2002; Meder, Gerstenberg, Hagmayer, & Waldmann, 2010). A person has an unbounded multitude of ways to try to make a golf shot, and exhaustive search over every possible strategy is not feasible. She must find a way to orient her evaluation toward strategies that are likely to be useful—strategies such as investing heavily in perfecting ball spin. This requires a preexisting representation of what kinds of strategies work “in general”—that is, averaging over a range of plausible circumstances (Hitchcock, 2012).

Token causal judgments help people acquire such a representations efficiently by employing a *sampling approximation*. That is, instead of acquiring such representations by exhaustive thought (explicitly considering all possible states of the system, and computing the effectiveness of an intervention in each one), our approach relies on past

experiences with the system—“samples” of that system, generated by the world and experienced by the person.

By analogy, humans and non-humans sometimes solve decision-making problems not by undertaking the computational effort of computing expected values from their representation of a complex system, but by approximating the values from their history of reward (Dolan & Dayan, 2013; Schultz, 2002). This works because the expected value of a future action can often be estimated as the average of its value in past episodes.

We apply a similar idea to causal judgment. For instance, the golfer can consider the specifics of many past shots, each of which represents a sample shot situation. By judging the causal status of a variable (e.g. “lots of spin”) in these episodes as she experiences them, the person can later obtain an estimate of the “average”, general role of a variable in the causal system. This, in turn, constitutes a useful starting point as she plans future shots.

Summary In sum, we propose that judgments of token causation help implement a sampling approximation of the overall effectiveness of an intervention. By repeatedly judging whether an event was a cause of an outcome in samples of a causal system, people estimate the average likelihood that intervening to produce the event would produce the outcome. This provides an efficient starting point for planning in future situations.

This paper proceeds as follows. We first review three prominent empirical patterns in token causal judgment, and summarize previous approaches to understanding them. We next present our approach, and show how, in a simple but common decision problem, causal judgments can help approximate the effectiveness of potential interventions by taking on a certain form. This form turns out to precisely reproduce the three empirical patterns. We then show that this form also provides a tight fit to people’s quantitative judgments in a recently reported experiment (Morris, Phillips, Gerstenberg, & Cushman, 2019). Finally, in the discussion, we consider how the principles of our approach may generalize to more complex decision problems, and ask how the novel predictions of our model might be tested.

3 The structure of token causal judgment

We begin by describing several basic structural features of token causal judgment established by prior research—in other words, features that predict which events people tend to select as “the cause” of some effect. Current theories do not provide a clear normative explanation for these patterns—in other words, they do not explain what goal is accomplished by selecting token causes in this manner. The algorithm we provide predicts each of these patterns and also explains their functional importance; thus, they constitute an important empirical basis for our proposal.

3.1 Conjunctive causal structures

Notably, the literature on token causal judgment has identified strikingly different patterns of judgment—indeed, often diametrically opposed patterns of judgments—that arise in the case of two different kinds of causal structure: “conjunctive” and “disjunctive” structures (Icard, Kominsky, & Knobe, 2017). In a conjunctive structure the effect of interest occurs only when *all* of the right causes are present—in other words, it is triggered by the conjunction of necessary elements. For instance, perhaps a safe is unlocked only if all three keys are turned at once. In contrast, in a disjunctive structure the effect of interest occurs when *any* of the right causes are present—in other words, it is triggered by the disjunction of sufficient elements. For instance, a different safe might be unlocked if any one of three keys is turned at all.

For several reasons we only focus on the structure of conjunctive causal structures. First, these patterns are consistent and well-understood. Second, they are also captured parsimoniously by our model. In fact, as we show in a later section, our model precisely tracks people’s causal judgments in a conjunctive system across systematic manipulations of the system’s free parameters.

These patterns, however, change significantly in disjunctive causal structures in which outcomes are causally overdetermined (or in which one cause preempts another, later-arriving cause).³ The basic version of our model struggles to explain judgments in

---

³In these cases, there are multiple events which are individually sufficient to produce an outcome—e.g. a prisoner is shot twice by two soldiers, and each shot would have been individually sufficient to kill him. A large amount of work has been devoted to understanding these kinds of cases; see Gerstenberg,

cases of overdetermination/preemption, and we revisit this issue in detail in the General Discussion.

Our assumption is that we can make some progress towards understanding causal judgment by introducing a model that addresses both its structure and function, even if that model is limited to conjunctive structures. Further work will be required to understand how the animating principles of this model can be adequately extended to disjunctive structures, or why they might fail to be extended.

More generally, we assume that our model will not provide a perfect fit to all known aspects of causal judgment. This is a natural consequence of its goal, which is not only to mirror the observed contours of human causal judgment, but also to provide a normative justification for those contours. In other words, the model is designed to show that human causal judgment is structured very much like an efficient solution to an important problem (i.e. Marr’s (1982) computational level of analysis.) It is unlikely that the exact algorithm we derive is implemented in the human brain. Rather, our key claim is that the functional design of the model we derive matches an important part of the functional design of causal judgment. This claim is motivated, in part, by the fact that the model reproduces the following three patterns in conjunctive structures: necessity, abnormal inflation, and supersession. We take these up in turn.

Feature 1: Necessity The first pattern in token causal judgment is that selected causes tend to be things that were *necessary* to produce the effect. Put simply, if the cause hadn’t happened in the specific relevant situation, the effect wouldn’t have either. For instance, if the camper hadn’t dropped a match, the forest would not have burned down; if the golfer hadn’t added spin, the shot would not have gone in; and so on. The fact that the camper’s match and the ball’s spin were necessary for their outcomes is a key determinant of their perceived causal status (Gerstenberg, Goodman, Lagnado, & Tenenbaum, 2014; Halpern & Pearl, 2005; Hart & Honoré, 1959/1985; Kahneman & Tversky, 1982; Lewis, 1973; Mackie, 1974; Wells & Gavanski, 1989; Woodward, 2003). In contrast, the camper’s decision to wear a blue shirt was not

---

Goodman, Lagnado, and Tenenbaum (2015); Hall (2004); Halpern (2016); Halpern and Pearl (2005); Hitchcock (2009); Lagnado, Gerstenberg, and Zultan (2013); Lewis (1973); Sartorio (2012).

necessary for the fire, and so it is not considered a cause.

Necessity plays a central role in theories of causation across diverse fields, from philosophy (Mackie, 1974) to law (Hart & Honoré, 1959/1985) and computer science (Chockler, Halpern, & Kupferman, 2008). Its role in causal judgment is also supported by a wealth of psychological evidence (Gerstenberg, Goodman, Lagnado, & Tenenbaum, 2012; Gerstenberg et al., 2015; Gerstenberg, Peterson, Goodman, Lagnado, & Tenenbaum, 2017; Kahneman & Tversky, 1982; Wells & Gavanski, 1989). In sum, there is widespread agreement that necessity, in some form, is important for judgments of token causation (Icard et al., 2017). (Of course, there is also widespread agreement that a simple necessity criterion is insufficient for distinguishing between causes and non-causes, as demonstrated by cases of overdetermination and preemption. We return to this issue in the General Discussion, and note how our model might be extended.)

While necessity helps distinguish causes (e.g. the match) from variables that are entirely noncausal (e.g. the color of the camper’s shirt), it does not tell us which variables will be considered more or less causal—a central issue for our account. The next two findings concern this issue.

Feature 2: *Abnormal inflation* The second pattern is that we tend to favor *abnormal* events as the token causes of an outcome (Fig. 1). Here, we focus on statistical abnormality: When multiple variables are necessary for an outcome, people prefer rare ones over common ones (Halpern & Hitchcock, 2015; Hart & Honoré, 1959/1985; Hesslow, 1988; Hilton & Slugoski, 1986; Hitchcock & Knobe, 2009; Kahneman & Miller, 1986; Willemsen & Kirfel, 2018).⁴ For instance, oxygen is common, but lit matches are rare; hence, people tend to consider the match more causal than the oxygen, even if both were necessary to make a fire burn. Similarly, golfers commonly make contact with the golf ball, but rarely impart the ideal amount of spin; hence, we are likely to consider the spin on a ball, rather than the mere fact the club hit it, as “the” cause of a difficult hole-in-one.

---

⁴People’s causal judgments also depend on *prescriptive* abnormality; they prefer variables that violate prescriptive or moral norms (Alicke, 2000; Hitchcock & Knobe, 2009) over ones that don’t. We return to prescriptive abnormality in the General Discussion.

Figure 1. (A) A conjunctive causal structure, where two variables are each necessary (and jointly sufficient) to produce an outcome. We illustrate these effects with an example similar to those from Kominsky et al. (2015) and Icard et al. (2017): A building’s climate control system activates if and only if two employees, Billy and Suzy, arrive at 9AM. (B) Results from Icard et al. (2017) demonstrating the “abnormal inflation” effect. Icard et al. presented participants with this causal structure, and qualitatively manipulated the prior probability of one of the variables. People judge Billy to be more of the cause of the climate system’s activation when he rarely arrives at 9AM. (C) Results from Kominsky et al. (2015) demonstrating the “supersession” effect. Using the same causal structure, Kominsky et al. manipulated the prior probability of the other variable. People judge Billy *less* causal when *Suzy* rarely arrives at 9AM. All results reproduced with permission of the authors. Error bars are ± 1 SEM.

The idea that causes tend to be abnormal is widespread. It plays an important role in numerous philosophical and legal theories of causation (Collingwood, 2014; Hart & Honoré, 1959/1985; Mackie, 1974), and has begun to be incorporated into formal models of token causation (Halpern & Hitchcock, 2015; Icard et al., 2017). And, like necessity, it is supported by a wealth of psychological evidence (Hilton & Slugoski, 1986; Hitchcock & Knobe, 2009; Kahneman & Miller, 1986; Kominsky, Phillips, Gerstenberg, Lagnado, & Knobe, 2015).

Feature 3: Supersession More recent research has identified a third signature pattern of token causal judgments: the perceived causal status of a given variable depends on the abnormality of *other* variables in the causal system (Kominsky et al., 2015, cf. McGill & Tenbrunsel, 2000). In a causal system where two variables are both needed to produce an outcome, people believe that one variable is less of a cause when the other variable is rare (Fig. 1C). For instance, the spin the golfer put on the ball seems less of the cause of her hole-in-one when a serendipitous gust of wind – a rare

event – drove it into the hole at the last second. Kominsky et al. (2015) labeled this effect “causal supersession”, capturing the idea that causal attribution to one event can be “superceded” by causal attribution to a less normal one (see also Icard et al., 2017; McGill & Tenbrunsel, 2000; Morris et al., 2019; Samland & Waldmann, 2016). Notably, just like abnormal inflation, this effect occurs when people already know the underlying causal structure, and when there are no inherent asymmetries in the “causal power” of the antecedent variables.

Summary In sum, when asked what caused an outcome, people tend to favor variables that were necessary for the outcome, variables that were statistically rare (“abnormal inflation”), and variables whose counterparts were *not* statistically rare (“supersession”). Necessity is one key way that people distinguish between causes and non-causes; abnormal inflation and supersession are two key ways that people select between identified causes.

In order to relate our discussion to other literatures, two points must be emphasized. First, these two processes—judging whether a variable was a cause at all, and judging which of the causes were “more causal” than others—are sometimes considered distinct, with the former at the heart of theories of token causation and the latter an auxiliary process of causal selection. Here, we treat them as a single integrated process that assigns a graded causal rating to variables. That rating can be zero (the color of the camper’s shirt), small (the oxygen), or large (the match) – but, in all cases, it is derived from a common process. In this way, our model aims to capture a functional motivation simultaneously underlying both basic token causal judgment and causal selection.

Second, abnormal inflation and supersession occur even when people already know the underlying causal structure. In other words, the patterns of causal selection that we are concerned with are not a result of inferences about causal structure due simply to sparse data. Moreover, people exhibit causal selection patterns even when there is no apparent asymmetry between the underlying “causal power” of the variables in that system. This is important because many models assume that the focal outcome is

produced by a “noisy OR” combination of two variables with different weights, and that asymmetries in causal attribution result from an inference about those weights (Cheng & Novick, 2005; Lu, Yuille, Liljeholm, Cheng, & Holyoak, 2008; Stephan & Waldmann, 2018). While surely part of the story, such models have difficulty explaining patterns of causal selection when several variables are equally necessary for an outcome (Icard et al., 2017). This latter case is the focus of our model.

4 The function of token causal judgment

Our goal is to show how an algorithm for causal judgment shaped by the *function* of supporting efficient planning also naturally accounts for the *structure* of causal judgment described above. In this section we provide a qualitative overview of the relationship between structure and function in causal judgment, and contrast our approach with other current theories.

4.1 Causal judgments identify effective interventions

We propose that the function of causal judgment is to assist in efficient planning. By planning, we have in mind a particular common kind of decision problem: We have a goal in mind, and our task is to find a way to accomplish that goal. It is easy to construe this decision problem in the language of a causal model: We wish to find a “cause” that will successfully bring about some desired “effect” in our present circumstances. Let’s return to the example of a golfer who wishes to make a hole-in-one. In order to bring about this “effect” she must assemble just the right set of causal conditions: The proper velocity of her swing, the proper angle, the proper follow-through, and so forth. (Construed as such, this is a conjunctive causal system: Making a hole-in-one depends on the appropriate conjunction of necessary causes.) In these kinds of decision problems, planning can profitably draw upon our causal knowledge.

It is natural to suppose, then, that our causal judgment are designed in part to support efficient planning. By judging what *has* caused an effect in many past episodes, we can efficiently discover what *will* cause that effect in the future, when we desire it.

Three key elements are required to establish this connection from past judgment

to future planning. First, we note that for every possible world in which *the effect and the cause are present* and *removing the cause would have removed the effect*, there exists a corresponding possible world in which *the effect and the cause are absent* and *adding the cause would add the effect*. Thus, judgments of “actual necessity” (the effect depended on the cause, in some specific past case) are logically related to useful goal-directed plans (introducing the cause would make the effect happen, in the future).

Of course, the logical correspondence between these worlds depends on holding everything else constant. In reality, however, circumstances change: The spin that landed a hole-in-one on a windy day may not be appropriate on a still day. Thus, we presume that the agent operates not with a detailed accounting of what-causes-what in each conceivable circumstance (“on the Pebble Beach course, hole 7, winds 3mph NNW...”), but rather generalizing across relevantly similar circumstances. Thus, she will average causal judgments across many past golf shots that differ in their particulars in order to arrive at a generic representation of what causes a hole-in-one “in general”. (In other words, she will approximate the optimal generic representation by employing a “sampling approximation” based on prior cases.) Although this generic representation cannot provide a precise answer to the question of what to do in any particular future circumstance, it can guide hypothesis generation, pointing a person towards potential candidate causes that have a high likelihood of being useful for bringing about the effect.

Generalizing across cases, however, means that we want not just “logical correspondence” between possible worlds in the past and possible worlds in the future but, in addition, “statistical correspondence”. That is, because we are generalizing across many possible worlds, we want to know just how often the particular worlds arise in which one cause or another is most useful. This problem is acute because the *more* often some cause was necessary to produce an effect in the past, the *less* often intervening on it will be to cause that effect in the future. For instance, consider the role of oxygen in a forest fire. For all past forest fires oxygen has been necessary; yet, it is almost never useful to consider the intervention “make oxygen” when one wants to *start* a forest fire. This is because oxygen is common—thus, reliably a causal force in

the past, and rarely the *missing* causal element in the future.

This establishes a functional explanation for “abnormal selection”, the tendency to rate rare causal conditions as strongly causal, while discounting common causal conditions as only weakly causal. More precisely: If one wishes to identify the (absent) causal conditions that will make a desired effect occur, and one wishes to do this by exploiting the logical connection to past circumstances where putative (present) causes were “actually necessary” to bring about the desired effect, then one must perform a statistical correction that depends upon the rate at which the candidate cause was absent *versus* present. In the following section we formalize this statistical correction and show that it corresponds precisely to the phenomenon of abnormal selection.

The same formalization shows that “supersession” also falls naturally out of the the appropriate statistical correction. So far we have emphasized that judgments of the past and future differ in terms of the presence of the “cause” (e.g., oxygen). If oxygen was present in the past, it is likely present in the future—thus, what was most often necessary in the past is least often a useful point of intervention. Note, however, that the past and the future also differ in the presence of the *effect*. When one judges the cause of a forest fire, or of a hole-in-one, the effect is present. When one wishes to make a forest fire, or to shoot a hole-in-one, however, the effect is absent. (That is precisely why one desires to “make” it!) As we show in the follow section, the appropriate statistical correction for this corresponds precisely to the phenomenon of supersession.

4.2 Relation to existing interventionist approaches

This proposal is partially inspired by recent “interventionist” approaches to causation in philosophy and psychology, but also departs from these in some ways. Thus, it is useful to briefly situate the proposal in a broader context.

Theories of causal judgment often commit, implicitly or explicitly, to a claim about what constitutes a causal relation. “Interventionist” approaches claim (roughly) that a causal relation is constituted by the fact that intervening on the putative cause would, at least in some circumstances, produce a change in the effect (Halpern & Pearl, 2005; Hitchcock, 2001; Woodward, 2003). Because we posit that the function of causal

judgment is to guide future intervention, our proposal dovetails elegantly with this perspective. (Of course, not all theories of causation assume a constitutive role for intervention; for other influential approaches, see, e.g., Cheng 1997; Lewis 1979; Spellman 1997).

Interventionist semantics—in particular, the languages of structural equation modeling and causal Bayes nets—have formed the backbone of a wealth of descriptive theories of causal judgment (Sloman, 2005). One strand of theories has modeled how people infer latent weights in noisy-OR causal structures from sparse data, and how they use these weights to make token causal judgments (Cheng & Novick, 2005; Griffiths & Tenenbaum, 2005; Holyoak, Lee, & Lu, 2010; Stephan & Waldmann, 2016, 2018). Another strand has modeled how people distinguish between causes and non-causes once they already know the underlying structure, with a focus on cases of overdetermination and preemption (Halpern & Pearl, 2005; Hitchcock, 2001; Woodward, 2003). A third strand, closest to our approach, has modeled how people assign *graded* causal ratings once they already know the underlying structure, with a focus on patterns like abnormal inflation and supersession (Halpern & Hitchcock, 2015; Icard et al., 2017; Kominsky et al., 2015; Phillips, Luguri, & Knobe, 2015).

Interventionist semantics are widely used because they offer a precise formal language for defining and modeling causal structure. But, although it has not always been emphasized, an exciting property of the interventionist approach is that causal judgment can be connected to *action* (Hagmayer & Sloman, 2009; Hitchcock, 2017; Meder et al., 2010; Sloman & Hagmayer, 2006). An interventionist framework can be used, not just in service of learning or defining causal structures, but also in service of actually deciding what to do—planning our own actual interventions in the actual world. This functional perspective grounds our proposal.

We call our approach the “Sample-based Approximation Method for Predicting the Likelihood of Effectiveness”—or the SAMPLE approach, for short. The SAMPLE approach illuminates the role of token causation in approximate planning, and illustrates how decision-makers can identify effective interventions, even in complex

scenarios where precise computation is difficult.

In the introduction, we mentioned that several prior theories have emphasized the connection between the formal notion of intervention as constitutive of causal relations (Pearl, 2000) and the practical notion of intervention that is integral to planning and action. In particular, thinking of interventions as actions a person could take to alter her environment has been proffered as a natural explanation for the existence of causal selection—that is, why people consider some causes more causal than others. If, in general, causes are variables which can be intervened on to manipulate effects, then perhaps the causes foregrounded by causal selection are particularly *good* interventions—ones that would work well as actions across a variety of situations (Hitchcock, 2012; Lombrozo, 2010; Vasilyeva, Blanchard, & Lombrozo, 2016; Woodward, 2006).

Our SAMPLE proposal builds directly on this idea. As mentioned in the introduction, however, there are two possible “orderings of judgments” implicit in this idea (Figure 2). In principle, people could *first* identify good interventions (i.e., solving the planning problem) and *then* grade their causal judgments to match them. But this proposal runs into two difficulties. First, identifying “good interventions” is an extremely nontrivial problem. Second, it is not clear what purpose is served by constraining the causal judgment of past events by the solution to this problem. As Hitchcock (2017) puts it, there is an apparent mismatch between the orientation of causal judgment (towards the past) and planning (towards the future). Token causation is “backward looking” (Hitchcock, 2017); it is anchored in the idiosyncratic details of the past event whose cause is being identified. In contrast, knowing which interventions tend to bring about an event in general is a forward-looking problem—the problem of decision-making. What purpose, then, is served by judging *specific, past* causes against the yardstick of *general, future* utility (Lombrozo, 2006)?

Thus, the SAMPLE model proposes the opposite relationship between these variables (Figure 2). It proposes that we make causal judgments of past events that take a specific form—a dependence on actual necessity, abnormal inflation, and

Two interpretations of the interventionist approach to causal selection

Figure 2. Schematic depiction of two possible relationships between causal selection and effective interventions. C represents the potential cause (e.g. the match); E represents its effect (e.g. fire).

suppression—because these aggregate efficiently to solve the difficult problem identifying “good interventions”.

Our next challenge is to specify the algorithm that accomplishes this goal. In the remainder of the paper, we formalize this approach in the language of structural equation modeling and, focusing on a simple but common decision problem, we show what form token causal judgments would have to take in order to accumulate evidence about effective interventions.

5 Formal model of effective intervention and token causation

In this section, we take the high-level proposal developed in the last section – that causal judgments accumulate information about the general effectiveness of future interventions – and develop it formally in a specific decision problem.

Imagine that a person wants to bring about an outcome E , such as starting a fire, getting a paper accepted, or making a shot in golf. The person has a model of the causal system in which E is situated.⁵ Moreover, the person can intervene on the system by

⁵Throughout our proposal, we will assume that the person making judgments and decisions already has a causal model of her environment. This assumption allows us to sideline the difficult problem of

turning on a single binary variable, C_i , from a set of candidate variables $C_1, \dots, C_K$. (By “turning on” a variable, we mean setting its value to 1.) For instance, the golfer trying to make a shot could focus on adding spin to the ball, changing her swing speed, using a different club, improving her follow-through, and so on. Which one should she choose?⁶

Structural equation modeling (SEM) (Pearl, 2000) offers a formal language for describing causal systems, and thus allows us to precisely define the effectiveness of potential interventions. We therefore situate our model in the SEM framework. For interested readers we provide a very brief introduction to the framework, with specific examples relevant to our proposal, in Appendix *i*.

5.1 Defining intervention effectiveness

A decision-maker is interacting with a causal system with variables, W . One variable, $E \in W$, is a desired outcome. Our decision maker must pick one of k variables, $C_1, \dots, C_k \in W$, on which to intervene. Moreover, she knows something about the situation; she knows the value of some subset of the variables $W_{known} \subseteq W$.

In this setup, the optimal intervention is clear: The decision-maker should choose C_i that maximizes the post-intervention probability of $E = 1$ conditioned on $W_{known} = w_{known}$. Formally, she should choose C_i that maximizes $Prob(E_{C_i=1} = 1 \mid W_{known} = w_{known})$. (This notation, specific to the SEM framework, reads as “The probability that $E = 1$ given the hypothetical intervention ensuring $C_i = 1$, given $W_{known} = w_{known}$ ”. More generally, it assigns probabilities under counterfactual or hypothetical interventions, which are subscripted).

Call this the *situation-specific effectiveness* of intervention C_i . Concretely, if a

---

causal induction (how people infer causal structures from sparse data; Bramley, Dayan, Griffiths, & Lagnado, 2017; Bramley, Gerstenberg, Mayrhofer, & Lagnado, in press; Coenen, Rehder, & Gureckis, 2015; Griffiths & Tenenbaum, 2005, 2009; Meder et al., 2010; Rottman & Hastie, 2013; Stephan & Waldmann, 2016; Steyvers, Tenenbaum, Wagenmakers, & Blum, 2003), and focus exclusively on causal selection — why people privilege some causes, even when they know the structure of their environment. This model need not be explicit, complete, or accessible to conscious deliberation; the only requirement is that people can perform the computations described below (e.g. assessing necessity, computing base rates).

⁶Our analysis assumes that the potential causes/interventions are binary (or can be treated as such for the purposes of the decision maker). This assumption, shared with other formal treatments of causation (e.g. Icard et al., 2017), makes our analysis tractable, but may often be unrealistic; see Appendix *iii*. For simplicity of exposition, we will notate the two values each variable can take as 0 (“off” or absent) and 1 (“on” or present). This, however, is purely notational; 0 and 1 can be conceptually replaced with whatever variable settings are relevant in the current context.

golfer is trying to make a shot and knows that the hole is close but it’s windy, then the situation-specific effectiveness of adding spin would be

$Prob(\text{MakeShot}_{\text{AddSpin}=1} = 1 \mid \text{HoleClose} = 1, \text{Windy} = 1)$. Situation-specific

effectiveness is, in this framework, the most accurate measure of the effectiveness of an intervention.

Yet, there is a strong incentive for decision-makers to also develop representations of the *general* effectiveness of interventions. Concretely, it is useful for a golfer to know what, in general, are the best things to focus on when taking a stroke. This information is critical because it helps direct the decision-maker’s evaluative resources towards interventions that are likely to have high situation-specific effectiveness. Most real-world situations have an enormous number of interventions to evaluate. (Think of all the ways a golfer could alter her stroke, or, e.g., all the ways a chef could improve a dish.) Without any idea of which interventions are generally effective, a person would be stuck evaluating all of them, or sampling a random subset to evaluate, and could waste an enormous amount of time and energy evaluating interventions that are not likely to work.

General effectiveness as a decision prior For this reason, decision-makers likely develop a representation of the “general” effectiveness of interventions, averaging across situation-specific details (more formally, marginalizing over all variables $V \neq E \in W$). Other authors have proposed a similar idea—that a useful representation of good interventions should be “robust” or “portable”, rather than being tethered to a highly specific state of affairs (Franklin & Frank, 2018; Hitchcock, 2012; Lombrozo, 2010; Vasilyeva et al., 2016). For instance, a person might know that lighting a match is, in most situations, an effective way to start a fire—even if there are some rare situations where it is ineffective (e.g., in strong winds). This representation would serve as a kind of prior on the true effectiveness of an intervention in a particular situation, which would be useful for two reasons. If someone did not have time to compute the effectiveness of each potential intervention conditioned on their situation-specific knowledge, they could quickly choose an intervention which was generally effective. And

if someone did have time, the “general effectiveness prior” would point to good places to start the more difficult computation. (Evaluating the situation-specific effectiveness of lighting a match is likely a better investment of computational resources than evaluating the effectiveness of pumping oxygen into the air.) Indeed, recent AI approaches to decision-making in large choice spaces often crucially depend on having good priors that guide choice evaluation, as in Monte Carlo tree search (e.g. Silver et al., 2016).

This generalized representation doesn’t have to completely abandon *all* situation-specific knowledge, however. Across most settings where a decision-maker is choosing interventions to produce an outcome, there is at least one constant: The outcome does not currently obtain. When someone is trying to start a fire, they typically believe there is not currently a fire; when a student hires an SAT tutor, she typically believes that her SAT scores are not high enough for college acceptance. (If the outcome already obtained, then the person would probably not be employing decision-making machinery on how to obtain it.) In other words, if we’re trying to identify generally effective interventions in contexts where a person is trying to bring about $E = 1$, we can safely condition on $E = 0$.⁷

This leads us to our definition of an effective intervention. The general effectiveness of intervening on C for bringing about E is $Prob(E_{C=1} = 1 \mid E = 0)$.⁸

Definition 1. The **general effectiveness** of intervening on C to produce E , denoted $GE(C \rightarrow E)$, is $Prob(E_{C=1} = 1 \mid E = 0)$.

To reiterate, we propose that judgments about whether C was the token cause of E help approximate the general effectiveness $GE(C \rightarrow E)$, and that $GE(C \rightarrow E)$ can

---

⁷There are, of course, many situations in which a decision-maker is unsure whether the outcome currently obtains. For instance, a student might hire an SAT tutor even though she thinks her application is already good enough — just to be sure. In these situations, it is still useful for a general effectiveness measure to condition on $E = 0$. So long as the decision-maker knows that an intervention won’t *decrease* the probability of the outcome (the common “monotonicity” assumption (Pearl, 1999)), then the choice of intervention is only important when the outcome doesn’t already obtain; the student’s decision whether to hire a writing tutor or sign up for debate team only matters if her application is currently insufficient. Hence, in these situations, a measure of general effectiveness should still condition on the absence of the outcome. (This analysis becomes more complex if there are intervention costs; see Appendix *ii* for a discussion.)

⁸In previous research, this value has been referred to as the “probability of enablement” (Pearl, 1999). We do not use this name to avoid confusion with the verb “enable”.

help a person make efficient decisions in future settings when she desires to bring about E . For instance, if the golfer has made causal judgments which imply that adding spin is generally an effective intervention, then, in a time crunch, she can quickly choose to do that; or, if she has time, she can focus on computing the situation-specific effectiveness of adding spin, and avoid wasting resources on less generally useful strategies.

General effectiveness, however, can be hard to compute explicitly. It requires evaluating all possible joint settings of the unknown exogenous variables—in plain English, computing the probability of all conceivable states of affairs in which the intervention might be applied. For instance, for the golfer to explicitly compute the general effectiveness of adding spin, she’d have to marginalize over situations when it’s windy or calm, rainy or sunny; when the hole is distant or close; when there’s obstacles or not; when she’s using irons or putters; and so on. (Formally, she would have to find all joint settings of the unknown variables under which, after adding spin, the shot would go in, and then sum their probabilities; see Fig. 6B in Appendix.) In a system with 10 binary exogenous variables, she would have to, in the worst-case scenario, evaluate over a thousand possible scenarios. It may be beneficial to avoid such computation.

5.2 Token causal judgments can approximate general effectiveness from experience

An alternative to explicitly computing the general effectiveness of an intervention would be to approximate it from experience. After all, each specific situation a person has experienced is a *sample* from the general distribution over situations. Put somewhat differently, each time a person observes a causal system, they are observing a snapshot of the variable settings. If these snapshots are randomly sampled, they can replace the computationally taxing process of enumerating over states of affairs. For instance, instead of evaluating all possible future shot situations, the golfer can rely on past shots as samples.

Crucially, however, a person must *do* something with the samples they observe; they must perform some computation on each observation. We propose that judgments

of token causation are the output of this computation, and when averaged across many samples, they approximate $GE(C \rightarrow E)$. Our aim is to specify a plausibly efficient algorithm that renders causal judgments implementing this approximation.

Suppose a decision-maker has a series of observations of the system $u_1, \dots, u_k$, where each u_i is a joint setting over the system’s variables (e.g. in u_7 , there was an arsonist, the ground was dry, and there was oxygen in the air). Suppose that she restricts her analysis to members of this set where $E = 1$ —i.e. when the outcome of interest is present. (Assume these experiences were randomly drawn from the system’s prior distribution, conditioned on $E = 1$.) Then, to approximate the general effectiveness of C for E , she should judge (1) whether $C = 1$ was necessary to produce the outcome in each case, (2) if it was, apply a correction term that weights rare causes and common effects more heavily, and (3) average those judgments together.⁹

We propose that a key function of token causal judgments is to implement the first two steps of this procedure. To judge the extent to which C was a cause of E in a particular sample u_i , a person determines whether C was necessary for E in u_i , and, if it was, weight it according to (among other things) the rarity of C . These judgments then assist in decision-making, because the person can simply maintain a running average of them to estimate $GE(C \rightarrow E)$.

Formally, we will use the notation $TC_{SAMPLE}(C \rightarrow E, u_i)$ to denote the extent to which C was a token cause of E in observation u_i (according to our SAMPLE model). Definition 2 presents our analysis of token causation.

Definition 2. On our model, the extent to which C was a **token cause** of E in a situation with exogenous settings u_i is:

$$TC_{SAMPLE}(C \rightarrow E, u_i) = \begin{cases} \frac{Prob(C=0) Prob(E=1)}{Prob(C=1) Prob(E=0)} & \text{if } C \text{ was necessary for } E \text{ in } u_i \\ 0 & \text{otherwise} \end{cases}$$

If causal judgments are guided by this formula, then the general effectiveness of C

---

⁹This includes situations where C is absent — i.e. $C = 0$. In these cases, we assume that C is automatically “unnecessary” and assigned a causal rating of zero.

for E can be estimated by simply averaging the causation judgments from each situation (see Proposition 1).

Proposition 1. If $u_1, \dots, u_k$ are randomly sampled joint settings of the exogenous variables (conditioned on $E = 1$), then, for large k :

$$GE(C \rightarrow E) \approx \frac{1}{k} \sum_{i=1}^k TC_{SAMPLE}(C \rightarrow E, u_i)$$

A proof of this proposition is provided in section *iv* of the Appendix.

What are the boundary conditions for this approximation method?

Importantly, this proposition holds true for all types of causal structures, conjunctive and disjunctive (and arbitrary combinations thereof). Earlier, we described how our model provides a compelling *descriptive* fit to judgments in conjunctive structures. This is true. But the *functional* aspect of our model—the idea that causal judgments can average out to an estimate of general intervention effectiveness—is not limited to conjunctive structures. It is widely applicable.

There is only one condition, the “no-confounding” condition, which must hold for this approximation to work. This provides an important functional boundary condition for our model. We describe this condition later on.

6 Model exposition: Unpacking the definition of actual cause

Having presented a formal model of the SAMPLE approach, our next goal is to discuss its main features in more detail.

We begin by illustrating how it works in a concrete case. Consider again the golfer. If her causal judgments of past shots align with Definition 2, our suggestion is that she will judge spin the cause of a particular shot going in to the extent that:

1. The shot’s success depended on the ball’s spin (necessity), i.e. C necessary for E in u_i ,
2. She rarely adds that much spin, i.e. $\frac{Prob(C=0)}{Prob(C=1)} \gg 0$, and

3. Shots are commonly successful, i.e. $\frac{Prob(E=1)}{Prob(E=0)} \gg 0$ (which, as described below, predicts supersession in conjunctive structures).

Having made these causal judgments in the cases of past shots, she can then approximate the general effectiveness of adding spin by simply averaging her past causal judgments. Such an algorithm is valuable because it allows her to estimate the general effectiveness of interventions in a tractable manner. She can then use that knowledge to guide her search for effective interventions in planning, either by choosing an intervention with high general value, or by devoting situation-specific planning especially towards interventions with high general value.

Why does this approximation work? We want to know the general effectiveness of getting from $E = 0$ to $E = 1$ by turning on C (i.e. $GE(C \rightarrow E)$). Clearly, this could be estimated by just computing the average effectiveness of C in past, randomly sampled episodes. One would simply ask: For what proportion of actual past episodes in which $E = 0$ did I successfully get $E = 1$ by turning on C ? (This is not a form of causal judgment, per se, but something more like reinforcement learning; Sutton & Barto, 1998.) This simple approach may not accumulate to a sufficiently accurate estimate when the opportunities to intervene in such situations are rare.

In the SAMPLE approach, token causation implements an alternate approximation—one that does not depend on trying actual interventions when an effect is absent, but that instead enables learning from counterfactual simulations over episodes where the effect is present. The existence of an episode where E was 0 and turning on C worked—call it episode u_i —implies a corresponding possible episode u'_i in which $E = 1$ and the agent *could* have turned C off in order to get $E = 0$. Causation judgments use the u'_i episodes, where turning C off would have turned E off, to learn about the u_i episodes, where turning C on would turn E on. It is as if we are mentally flipping a light switch on and off, watching a specific room go light and dark. If flipping C on brings a specific room u_i from darkness to light, then flipping C off must bring a corresponding room u'_i from light to darkness. Establishing the correspondence of such rooms—or states u_i and u'_i —is the job of token causal judgments.

But we are not just trying to establish the existence of such rooms—we want to know their probabilities of occurrence, which tells us how often C is a good point of intervention. (As described in section 4, we want not only “logistical correspondence”, but also “statistical correspondence”.) Of course, it will not do to simply compute the probability that episodes like u'_i naturally arise, and then treat this as an estimate of the probability that episodes like u_i naturally arise. Although the *existence* of any u'_i implies the *existence* of a corresponding u_i , their relative *probabilities of occurrence* will usually differ. Thus, in order to compute the probability of episodes like u_i from samples of episodes like u'_i , we need to perform a statistical correction for their relative probabilities of occurrence. This is accomplished with the correction term $\frac{\text{Prob}(C=0) \text{Prob}(E=1)}{\text{Prob}(C=1) \text{Prob}(E=0)}$. Having done this, we can accumulate evidence from episodes in which the effect was present, and could be eliminated by turning C off, in order to estimate the probability of episodes in which the effect was absent, and could be produced by turning C on.

In the remainder of this section we further develop the intuition for why causal judgments of this form can approximate the general effectiveness of interventions, and how they provide a normative explanation for the empirical patterns described above. We also describe the condition—the no-confounding assumption—that must hold for this approximation to work. The core of our proposal is in Definition 2—the form that token causal judgments take to approximate intervention effectiveness. There are two parts to the definition: the necessity criterion (whether turning off the potential cause would have turned off the outcome), and the correction term. We take up each in turn.

6.1 Token causes should be necessary

In the SAMPLE approach, if token causation were designed to approximate general intervention effectiveness, then the first step in determining causes would be to judge whether “turning off” the potential cause—i.e. intervening to set $C = 0$ —would have turned off the outcome. For instance, to determine whether the lit match caused the forest fire, a person would judge whether intervening to prevent the match from being lit would have prevented the forest fire.

This corresponds to the standard necessity criterion: Judging whether turning off C would have turned off E is a way of assessing whether C was necessary for E . What is promising about this connection is that starting from normative assumptions about general effectiveness, we can capture the ubiquitous finding that necessity is central to causal judgment.

At the same time, this result might seem puzzling. In our framework, the general effectiveness of an intervention is about whether turning a variable *on* (e.g. lighting a match) would produce an outcome when the outcome is absent—but it's being estimated by considering whether turning a variable *off* (e.g. snuffing out a match) would have prevented an outcome when the outcome is present. Put differently, the definition of $GE(C \rightarrow E)$ involves setting $C = 1$ when $E = 0$, but the definition of $TC_{SAMPLE}(C \rightarrow E, u_i)$ requires people to imagine setting $C = 0$ when $E = 1$. Why?

The first thing to note is that our approximation algorithm is not meant to be the only solution to the problem of identifying effective interventions; the general effectiveness of an intervention could be estimated in other ways. People *could* estimate $GE(C \rightarrow E)$ more directly by determining whether turning on the variable (e.g. lighting a match) would produce the outcome in situations where the outcome is absent. In other words, instead of relying on judgments of necessity, people could rely on judgments of sufficiency. In fact, at times, people probably do employ this alternative approximation method. Any time a person considers whether turning on a variable would produce an outcome in a given situation, they can use that computation to inform their future decision-making more generally. Sufficiency judgments likely play a role in assessing the effectiveness of interventions, but we would typically describe this as a form of planning, rather than causal judgment.¹⁰

While it is intuitive that sufficiency judgments may play a role in decision making,

---

¹⁰This alternative approximation algorithm would be simple. Suppose we have observations of the system $z_1, \dots, z_k$, which are randomly sampled from the prior over exogenous variables, conditioned on $E = 0$. In other words, we observed a series of instances where there was no fire, or when no paper was accepted. To estimate $GE(C \rightarrow E)$, we could simply determine whether $C = 1$ (e.g. lighting a match) would have been sufficient to produce $E = 1$ in each case z_i , and average those judgments together. That is: $GE(C \rightarrow E) \approx \frac{1}{k} \sum_{i=1}^k I(E_{C=1} = 1 \mid Z = z_i)$, where I is a function that equals 1 if its argument is true, and 0 otherwise. This alternative approximation could then be averaged with the results of the necessity approximation to obtain a better estimate of $GE(C \rightarrow E)$.

the key insight of our proposal is that necessity judgments—of the type employed in causal judgment—can also play a role, and indeed potentially a very large one. As long as the causal system satisfies certain conditions described below, and after applying the correction term specified in our algorithm, judging whether turning off a variable would have turned off an outcome provides an unbiased estimate of that variable’s ability to produce the outcome.

Necessity, sufficiency, and the no confounding assumption As we have emphasized, the conversion from necessity to sufficiency works because any set of background factors (i.e. exogenous variable settings) that makes a variable necessary (when it is on) is identical to the set that makes a variable sufficient (when it is off). But this relationship—between whether a variable is necessary, and whether it would be sufficient—depends on a key assumption: The probability of the relevant background factors must not depend on the setting of the variable. Concretely, judging when the presence of spin was necessary provides an unbiased estimate of when adding spin would be sufficient, so long as the probability of the relevant background factors (e.g. the wind, the type of club, etc.) does not depend on the ball’s current level of spin.¹¹

The commonly employed “no-confounding condition” guarantees that the situations which make a variable necessary when it is on will be equally probable when the variable is off. The technical definition, modified for our context, is as follows:

Definition 3. The **no-confounding condition** holds for a potential cause/effect pair (C, E) if and only if:

$$Prob(E_{C=1} = 1, E_{C=0} = 0 \mid C = 0) = Prob(E_{C=1} = 1, E_{C=0} = 0 \mid C = 1)$$

In other words, the effect on the outcome E of intervening on the potential cause C must be independent of whether C is on or off. (Equivalently, the variable settings that set the backdrop for intervening on C are equally probable when C is on or off.)

This technical criterion can be intuitively represented by a graphical criterion. The

---

¹¹A variable is “relevant” here if it affects the post-intervention probability of the outcome. Of course, a decision-maker does not have to *judge* which background variables are relevant; rather, adherence to the no-confounding assumption is a property of environments that make them amenable to our analysis.

no-confounding assumption is guaranteed to hold if C and E have no common ancestor—if there is no upstream variable that influences both the potential cause and its effect.

In light of the previous discussion, it is easy to see why the no-confounding assumption supports a conversion from necessity to sufficiency. It guarantees that the background factors which make a variable necessary when it is on are equally likely to occur when the variable is off. The only difference between the necessity-situations and sufficiency-situations are the settings of the focal variable C and its outcome E —and these get accounted for in the correction term $\frac{Prob(C=0) Prob(E=1)}{Prob(C=1) Prob(E=0)}$, which we discuss below. (For a technical description of the no-confounding assumption’s role in our model, see section *iv* in Appendix.)

How plausible is this assumption in real-world causal systems? Though rarely perfectly fulfilled, the no-confounding assumption is plausible enough to be widely adopted in practical fields like epidemiology and law (Pearl, 2000), and is a weaker assumption than those made in popular psychological models (Cheng, 1997; Luhmann & Ahn, 2005). We return to the role of the no-confounding assumption in the general discussion.

Why use necessity judgments? Judgments of actual necessity are not the only way to approximating intervention effectiveness. A sufficiency approximation (i.e., imagining $C = 1$ when $E = 0$) is an obvious alternative, and it has *prima face* appeal because one can directly aggregate specific sufficiency judgments into a general representation of intervention effectiveness without statistical correction.

There are, however, two reasons that the necessity approximation might be a particularly rich and efficient method for accumulating evidence about effective interventions. First, when people have a goal of obtaining some outcome $E = 1$, it is typically the case that the presence of the outcome is rare, compared to its absence. People observe many, many more instances with no fire than instances with fire. As a consequence, it would be quite difficult to judge sufficiency in all the instances where $E = 0$. (Imagine asking yourself, every moment of every day, which variables would be

currently sufficient to produce a fire.) If similar estimates can be obtained either way, then it is computationally economical to focus on the relatively few instances where $E = 1$.

Second, the structure of many important causal systems might make it easier to assess necessity than sufficiency. Often, when an outcome occurs, people find it easy to generate many variables that were necessary for it. In contrast, when an outcome doesn't occur, it is more difficult to generate variables that would have been sufficient to produce it. For instance, imagine that a basketball player makes a 3-pointer. It is easy for even a casual fan to think of ways he could have missed—if his arc had been a little lower, if the defending player had jumped a little higher, if his knee had given out, and so on. But, if a player is missing his 3-pointers, it may be difficult to describe exactly what he should change to sink them. (If it were just that easy, basketball coaches would be out of a job!)

Of course, there are also many cases where sufficiency is intuitively easier to assess than necessity. An important task for future research is to identify the characteristics of causal systems that make necessity or sufficiency more salient. For now, we note that necessity judgments seem salient in enough cases that decision makers would likely profit from a cognitive tool—token causation—dedicated to harvesting the information they provide.

In sum, by judging whether a variable was necessary for an outcome when the outcome was present, people can determine whether a variable will produce an outcome when it is absent. Critically, however, this approximation only works if people apply an appropriate correction term. We now turn to discussing this correction term, show why it is important, and show how it provides a surprisingly tight fit with the two other empirical patterns we noted at the outset: “abnormal inflation” and “supersession”.

6.2 The correction term

The second part of our definition of token cause (Definition 2) is the correction term: If C was necessary for E , then the extent to which it was a token cause is:

$$\frac{Prob(C = 0) Prob(E = 1)}{Prob(C = 1) Prob(E = 0)}$$

This term allows necessity judgments when $E = 1$ to accurately approximate the general effectiveness of potential interventions when $E = 0$.

There are three important features of this correction term. The first is that it makes token causal judgments graded, not all-or-nothing; in other words, one variable can be more of a cause than another. This accords with a wealth of empirical evidence that people, in fact, treat token causation as graded (Danks, 2017; Gerstenberg et al., 2015, 2017; Halpern & Hitchcock, 2015; Kominsky et al., 2015).

Token causes should be abnormal The second feature is that it weighs rare variables more heavily (via $\frac{Prob(C=0)}{Prob(C=1)}$). As discussed above, this also accords with a wealth of empirical evidence that people favor rare causes over common ones—the “abnormal inflation” effect (Hilton & Slugoski, 1986; Hitchcock & Knobe, 2009; Icard et al., 2017). There are two reasons that abnormal causes are weighted more heavily. The first, captured by the $Prob(C = 0)$ term in the numerator of Definition 2, follows from a basic fact: Turning a variable “on” cannot be an effective intervention in any situation in which it is already “on”. (For instance, supplying oxygen cannot be an effective way to start a fire if there is already plenty of oxygen, and yet no fire). Thus, in order for turning a variable on to be a good intervention in general, the variable must generally be “off”.¹²

Consider again a person trying to turn an outcome from off to on (e.g. “shot missed” to “shot made”), who can intervene to turn a single variable on. In this scenario, the only interventions that can make a difference are interventions on variables

¹²This exposition assumes that candidate interventions always involve turning a variable on, as opposed to off—i.e. introducing a factor which was absent. Of course, effects are sometimes produced by *removing* a factor — e.g. erasing a student’s criminal record may help them get into college. This possibility is easily incorporated into our framework; “on” and “off” can be reinterpreted, not as the presence or absence of a variable, but as the relevant settings of the variable in the current context.

that are *currently off*. If she’s already put the ideal spin on the ball, then intervening on the spin cannot help. Hence, a measure of intervention effectiveness should be weighted towards variables that are more likely to be off.

The second reason why abnormal causes are weighed more heavily is captured by the $Prob(C = 1)$ term in the denominator of Definition 2. This reason is a counterpart to the first reason. Recall that in our model (and in typical human judgments), causes must have been necessary for the outcome. For a variable to be necessary, it had to be “on”. A lit match could not have been necessary for the fire that occurred if there were no lit match. Hence, by always considering variables which were necessary, we end up, on average, biasing our token cause judgments towards variables that tend to be present rather than absent. Dividing by $Prob(C = 1)$ in each particular case—i.e. focusing on abnormal variables—corrects for this bias.

Causes should have common outcomes Finally, the correction term weighs variables more heavily when the outcome is more common (via $\frac{Prob(E=1)}{Prob(E=0)}$). The reason for this weighting is that necessity judgments and general intervention effectiveness are restricted to different types of situations. Our approximation assumes that people are only making necessity judgments in situations where the outcome happens ($E = 1$). However, the effectiveness of an intervention is, in our definition, measured in situations where the outcome doesn’t happen ($E = 0$).

The $\frac{Prob(E=1)}{Prob(E=0)}$ term corrects for this difference. It follows from Bayes’ rule that, when conditioning observations on some criterion (e.g. $E = 1$), you divide by the probability of that criterion. In some sense, then, people’s necessity judgments have been “divided” by $Prob(E = 1)$; multiplying by $Prob(E = 1)$ undoes that conditioning. Moreover, we want our measure of intervention effectiveness to be re-conditioned on $E = 0$; hence we divide the measure by $Prob(E = 0)$.

Importantly, the $\frac{Prob(E=1)}{Prob(E=0)}$ term provides a normative explanation for the supersession effect. Consider again the two-variable conjunctive case used in Icard et al. (2017) and Kominsky et al. (2015). An alarm goes off if and only if two employees arrive at 9 AM. The supersession effect is that people consider one employee (Billy) less

causal if the other employee (Suzy) *rarely* arrives at 9 AM. This follows from our model: If Suzy rarely arrives at 9 AM, then, all else being equal, the alarm will go off less often—and hence Billy should be considered less causal.

Notably, this explanation of supersession differs from prior accounts. Supersession has been explained by assuming that, in addition to requiring causes to be necessary, people also require them to be “robustly sufficient”—i.e. sufficient across a variety of circumstances (Kominsky et al., 2015; Woodward, 2006). If Suzy usually arrives at 9 AM, it makes the alarm more likely to go off (the explanation considered here); but it also makes Billy more likely to be sufficient to make the alarm go off (the robust sufficiency explanation). Either of these theories, then, could explain the supersession effect in the two-variable conjunctive case.¹³

6.3 Summary

We have argued that judging causes of past outcomes can be used to efficiently approximate the effectiveness of potential interventions, if those judgments follow three guidelines: Causes should be necessary, statistically rare, and linked to common outcomes. These guidelines offer a functional explanation for three prominent empirical patterns: necessity, abnormal inflation, and supersession. Next, we apply the model to recently reported data, and find that the model provides a tight quantitative fit to people’s judgments in a conjunctive causal structure.

7 Applying SAMPLE to fit quantitative patterns of judgment

We have demonstrated that the SAMPLE approach can account for qualitative effects like abnormal inflation and supersession. Here, we show that it also provides a tight fit to certain quantitative patterns of judgment as well. (All data and analysis code used in this section can be found at <https://github.com/adammorris/SAMPLE>.)

We recently reported an experiment which measured the form of people’s

---

¹³An attentive reader may wonder: What differs between our account of token causation, and the robust sufficiency account? It seems like “robustly sufficient” is just another way of describing the general effectiveness of an intervention. The key difference is that, on our model, people don’t explicitly employ a robust sufficiency criterion when making causal judgments. Rather, something akin to “sufficiency robustness” is the value which causal judgments are designed to approximate. (The robust sufficiency account is similar to the “intervention first” account described in Figure 2.).

judgments in a simple conjunctive causal structure (Morris et al., 2019). We presented people with a vignette like the one in Fig. 1A, where two variables were both necessary for an outcome. We created one hundred different conditions by systematically manipulating the prior probability of both variables. In each condition, we told people that both variables and the outcome occurred, and asked them to rate the extent to which they believed one of the two variables was “the cause” of the outcome. This design allowed us to quantitatively measure how people’s judgments changed as a function of the prior probability of the potential causes. (For details of the experiment, see Morris et al. (2019).)

7.1 SAMPLE’s predictions

The SAMPLE model makes predictions about the form of people’s judgments in this causal structure. To repeat our definition, according to the model, the extent to which C is a cause of E in situation u_i should be:

$$TC_{SAMPLE}(C \rightarrow E, u_i) = \begin{cases} \frac{Prob(C=0) Prob(E=1)}{Prob(C=1) Prob(E=0)} & \text{if } C \text{ was necessary for } E \text{ in } u_i \\ 0 & \text{otherwise} \end{cases}$$

Causal structure in experiment

$$\text{Outcome} = 1 \text{ iff Focal} = 1 \text{ AND Alternate} = 1$$

Figure 3. The causal structure used in the experiment of Morris et al. (2019).

Let **focal** represent the variable whose causal status people are asked to judge; let **alternate** represent the other variable; and let **outcome** represent the outcome (see

Fig. 3). (To make the following equations more readable, we will use $Prob(\mathbf{focal})$ and $Prob(\neg\mathbf{focal})$ as shorthand for $Prob(\mathbf{focal} = 1)$ and $Prob(\mathbf{focal} = 0)$ — i.e. to indicate the probabilities of the variable occurring and not occurring, respectively. We will do the same for $\mathbf{alternate}$.) According to the SAMPLE model, people’s causal judgments in this experiment should take the following form:

$$TC_{SAMPLE}(\mathbf{focal} \rightarrow \mathbf{outcome}) = \frac{Prob(\neg\mathbf{focal}) \cdot Prob(\mathbf{alternate})}{Prob(\neg\mathbf{focal}) \cdot Prob(\mathbf{alternate}) + Prob(\neg\mathbf{alternate})}$$

This formula follows from applying the definition above to the two-variable conjunctive causal structure. It has an intuitive interpretation: It is the probability that, given that the outcome did not occur, intervening to turn on the focal variable would make it occur. In other words, it is the general effectiveness of $\mathbf{focal}$ for producing $\mathbf{outcome}$ — $GE(\mathbf{focal} \rightarrow \mathbf{outcome})$.¹⁴

This causal structure, then, represents a special case for our model. In general, token causal judgments are predicted to average to, not equal, the effectiveness of interventions (i.e. $GE(C \rightarrow E) = \frac{1}{k} \sum_{i=1}^k TC(C \rightarrow E, u_i)$; see Figure 2). But when a variable is always necessary for an outcome (as in this case), our model predicts that causal judgments will be identical in each episode—and hence equal to the average (in this case, $GE(\mathbf{focal} \rightarrow \mathbf{outcome}) = TC_{SAMPLE}(\mathbf{focal} \rightarrow \mathbf{outcome})$). This case removes the influence of the necessity criterion (since it is always satisfied), and allows us to test whether people’s judgments align with the form predicted by the more novel half of our model: the correction term.

Unpacking the predictions The SAMPLE model’s predictions are shown in Figure 4. It is helpful to unpack certain qualitative aspects of these predictions. First, as the focal variable becomes rarer, people should think it is more causal. This is abnormal inflation. Second, as the alternate variable becomes rarer, people should think the focal variable was less causal. This is supersession.

The novel predictions of the SAMPLE model concern the overall, nonlinear

¹⁴To see this, consider that there are two ways that the outcome could be off: Either $\mathbf{focal}$ is off and $\mathbf{alternate}$ is on, or $\mathbf{alternate}$ is off. Only in the former case would setting $\mathbf{focal}$ to 1 produce $\mathbf{outcome} = 1$. TC_{SAMPLE} captures the probability of the former case.

shape of the responses (Fig. 4). For instance, our model predicts a relatively weak effect of $Prob(\text{focal})$ on causal judgment when the probability of **alternate** is close to (or equal to) 1—except for a steep drop in the causality of **focal** from $Prob(\text{focal}) = .9$ to $Prob(\text{focal}) = 1$. Similarly, our model predicts a relatively weak effect of $Prob(\text{alternate})$ when the probability of **focal** is close to 1—except for a steep increase in the causality of **focal** from $Prob(\text{alternate}) = .9$ to $Prob(\text{alternate}) = 1$. These predictions are unique to the SAMPLE model, and follow from the model’s functional rationale. As described above, in this simple causal structure, the SAMPLE model predicts that people’s causal judgments should equal $GE(\text{focal} \rightarrow \text{outcome})$ —the probability that, given that the outcome didn’t occur, intervening on **focal** would make it occur. When $Prob(\text{alternate})$ is close to 1, this probability is consistently very high; the outcome hasn’t happened, it’s probably because **focal** = 0 (and thus setting **focal** = 1 would be a successful intervention). The only exception is when $Prob(\text{focal})$ is also close to 1. Then, suddenly, it is no longer near-certain that **focal** is the missing ingredient, and the probability of **focal** starts mattering again. This explains why the SAMPLE model predicts a weak effect of $Prob(\text{focal})$ when $Prob(\text{alternate})$ is close to 1, except for a steep drop when $Prob(\text{focal})$ approaches 1 as well. (Similar logic also governs the effect of $Prob(\text{alternate})$ when $Prob(\text{focal})$ is close to 1.)

Normalized causal judgments We also present the predictions of a variant of our model, in which people judge the causality of **green** in comparison to its alternate, **alternate**, by normalizing their response for $TC_{SAMPLE}(\text{focal} \rightarrow \text{outcome})$ with respect to $TC_{SAMPLE}(\text{alternate} \rightarrow \text{outcome})$. We discuss the arguments for normalization in Appendix *v*. The idea that people normalize causal judgments is controversial (Gerstenberg et al., 2015; Kominsky et al., 2015; Lagnado et al., 2013). We do not claim it as a general rule. Nonetheless, the quantitative nature of our experiment allows us to test whether a normalized model better captures people’s judgments in this causal structure.

A natural way to accomplish normalization is to enter the causal ratings into a

Figure 4. The predicted response profile of the SAMPLe model, as a function of the prior probability of **focal** and **alternate**. The left graph shows the response profile in three dimensions; the right graph shows the same profile, collapsed into two dimensions. The SAMPLe model predicts abnormal inflation (as the probability of **focal** increases, **focal** becomes less causal) and supersession (as the probability of **alternate** increases, **focal** becomes more causal); but it predicts that these effects will be nonlinear.

logistic (or “softmax”) function:

$$TC_{SAMPLe-normalized}(\text{focal} \rightarrow \text{outcome}) = \frac{e^{\beta \cdot TC_{SAMPLe}(\text{focal} \rightarrow \text{outcome})}}{\sum_{\text{var} \in \{\text{focal}, \text{alternate}\}} e^{\beta \cdot TC_{SAMPLe}(\text{var} \rightarrow \text{outcome})}}$$

where β is a “gain” parameter that controls the strength of normalization. People’s causal judgments would then reflect this normalized value.

The predictions of this are shown in Figure 5B. (To keep the model parameter-free, we fix β at 1.) The normalized model predicts an almost identical response profile as the non-normalized model, with two notable differences. The first is that the range is restricted; instead of rating the causality of **focal** as zero when, e.g., $Prob(\text{focal}) = 1$, the normalized model rates it at around 0.3. The second is that the normalized model predicts a more consistent abnormal inflation effect, even when $Prob(\text{alternate})$ is very low. In other words, the normalized model predicts that people will continue to rate **focal** as more causal when it is rarer, even when **alternate** are almost never drawn (see the downward slope on the light-gray lines in Figure 5B). To preview our results, both of these distinctive features are observed in human data.¹⁵

¹⁵Our model, in both its unnormalized and normalized forms, does not technically assign causal rat-

Figure 5. The predicted response profiles of our model, as a function of the prior probability of the two antecedent variables. The x-axis depicts the probability of *focal*, the variable whose causal status is being judged; the line shade depicts the probability of selecting *alternate*; and the y-axis depicts the predicted (or empirical) causality rating for selecting *focal*. (A/B) The predictions of the unnormalized/normalized SAMPLE model; (C) People's average responses; (D/E) the relationship between the model predictions and people's average ratings in each condition.

7.2 Results

People’s average ratings from the experiment in Morris et al. (2019) are shown in Figure 5C, alongside the predictions of the SAMPLE model (both unnormalized and normalized). Our model provides a tight quantitative fit to people’s judgments.

Reanalyzing the data from Morris et al. revealed that our unnormalized model had a correlation of $r = .86$ ($t(98) = 16.6, p < .001$; Fig. 5D), and our normalized model had a correlation of $r = 0.87$ ($t(98) = 17.7, p < .001$; Fig. 5E). Both correlations were significantly higher than that of any model analyzed in Morris et al. (Williams’ t-test, $ts > 2.5, ps < .01$), and they did not significantly differ from each other ($t < 1, p = .42$).

As predicted by our model, people exhibit markedly reduced abnormal inflation and superseding when $Prob(\text{alternate}) = 1$ and $Prob(\text{focal}) = 1$, respectively. The one exception, as predicted by our model, is when the probabilities of **focal** and **alternate** are both close to 1. In that case, people show steep abnormal inflation and supersession.

Notably, when qualitatively comparing the normalized and unnormalized versions of our model, people’s ratings were most consistent with the normalized version. The absolute value of their ratings largely fell within the normalized model’s restricted range (around 0.3 to 0.7), and they exhibited consistent abnormal inflation, even when **alternate** was very rare. These results provide some evidence that, in this experiment, people are likely normalizing their responses. An interesting question for future research is when, and why, people normalize when reporting causal judgments.

8 General Discussion

What are token causation judgments for, and why do people consider some token causes “more causal” than others? We argue that a key role these judgments serve is to accumulate evidence about generally effective interventions in a computationally efficient manner. Our model provides a natural, functional account of three qualitative patterns widely observed in causal judgments in conjunctive structures: necessity,

---

ings in the situation where $Prob(\text{focal}) = Prob(\text{alternate}) = 1$. (In this case, the denominator in $TC_{SAMPLE}(\text{focal}, \text{outcome})$ is zero.) However, if people do not assign a probability of exactly 1—e.g. if they assign a probability of 0.999—to $Prob(\text{focal})$ and $Prob(\text{alternate})$, then the causal rating of green approaches 0.5. Thus, we assign a rating of 0.5 in this case.

abnormal inflation, and supersession. It also captures participants’ quantitative causal judgments in a recently reported experiment, across systematic manipulations of the free parameters in a conjunctive system (Morris et al., 2019). This match between the predictions of a normative model of decision-making and patterns in judgments of token causation, at least in one class of environments, suggest that they share some common functional design.

More generally, our model addresses a basic question about the nature of token causal judgment. As noted by Hitchcock (2017), an organism’s fate is determined by the future consequences of its actions. Decision-making must, in this sense, be “forward-looking”. Yet humans appear obsessed with inherently backward-looking ascriptions of causal responsibility (see also Woodward, 2011). Why?

The SAMPLE model bridges this gap by invoking the notion of a sampling approximation. Past observations of a causal system can act as samples which, when combined appropriately, enable a person to approximate the future effectiveness of potential interventions. Viewing causation as an implementation of this sampling approximation explains why a forward-looking enterprise like decision-making would be served by a backward-looking notion like token causation (Lombrozo, 2006).

Our remaining discussion begins with a review of the boundary conditions of our analysis. We consider our core assumptions, and how they limit the scope of our approach. Then, we describe the connections between the SAMPLE model and previous work, and highlight how our model may be more deeply integrated with previous approaches. Finally, we consider future directions and open questions.

8.1 Assumptions and limitations

Our analysis has several key boundary conditions and assumptions, and we begin by considering three especially fundamental ones: The no-confounding assumption, the dependence on necessity judgments, and the case of disjunctive causal structures.

Functional boundary condition 1: the no-confounding assumption

The first core assumption is that token causation primarily functions in causal structures that obey the “no-confounding” assumption: The effect on desired outcome

E from intervening on potential cause C does not depend on the current setting of C . Graphically, this assumption is guaranteed to hold if the cause and effect have no common ancestor; there are no “confounds” connecting them. This assumption is an indelible commitment of the SAMPLE approach to causation because it enables necessity judgments to inform judgments of intervention effectiveness; without it, necessity-based causation judgments cannot accumulate evidence about which interventions are likely to produce outcomes.

Fortunately, it is reasonable to suppose that causal judgments are at least partially predicated on the no-confounding assumption. The no-confounding assumption is met (or partially met) often enough to warrant its widespread adoption in fields like epidemiology, law, and econometrics (Pearl, 1999, 2000). Moreover, from a psychological perspective, the no-confounding assumption is weaker than the assumptions made in popular models like Power-PC (Cheng, 1997; Luhmann & Ahn, 2005), and people are anecdotally hesitant to make causal judgments when it is violated (Luhmann & Ahn, 2005).¹⁶ Finally, from a functional perspective, adopting the no-confounding assumption is the only way to extract information about intervention effectiveness from necessity judgments. Tian and Pearl (2000) show that, without the no-confounding assumption, the probability that an intervention successfully produces an outcome is not bounded by information about necessity. If people face many situations where they (a) must choose interventions to produce outcomes, and (b) can identify necessary variables from past experience, then a reasonable approach is to adopt the no-confounding assumption and extract whatever information can be gleaned from that past experience.¹⁷

Functional boundary condition 2: Ease of assessing necessity The no-confounding assumption is the only condition required for our approximation algorithm to work. There is, however, a second core assumption that motivates *why*

---

¹⁶An important caveat to these observations is that the no-confounding assumption is most often adopted in treatments of *type* causation, not token causation. The extent to which the assumption translates to cases of token causation has been less explored.

¹⁷Note that if the no-confounding assumption fails, necessity judgments will not systematically bias people *away* from effective interventions; they will simply be uninformative, and add noise to the estimate. Hence, a cost-benefit analysis might reasonably conclude that making the no-confounding assumption, and using necessity judgments, is worthwhile.

people would use this approximation: We assumed that people would be estimating intervention effectiveness in scenarios where necessity is easier to assess than sufficiency. In other words, we assumed that it would be easier to identify variables which were necessary for an outcome than to determine the probability that a variable would be sufficient for the outcome. This assumption helped motivate why people would use necessity judgments, rather than sufficiency judgments, to estimate the effectiveness of interventions—and hence why causal judgments would take the form that they do (Gerstenberg et al., 2014; Halpern & Pearl, 2005).

This assumption is clearly met in conjunctive situations, where many variables are necessary for an outcome (e.g. a fire occurs just in case there’s oxygen, and dry wood, and a lit match, and so on) and none are solely sufficient. But it likely holds in many other structures. There are many causal structures with disjunctive elements—“OR” parameterizations—in which necessity is still easier to identify and assess. Consider shooting a basketball. There is more than one way to make a successful shot—you can release the ball lower or higher; you can jump towards the basket, or away from it; etc.—suggesting that there are disjunctive elements in this causal system. Yet, it is relatively easy to identify variables which were necessary for a particular shot to go in, and relatively difficult to identify sets of variables that would be sufficient for a shot to go in. (For a difficult shot, any casual fan can identify necessary conditions for a successful shot, but only a skilled player can identify sufficient conditions for a shot-to-be-made). This pattern is likely replicated in any real-world structure where there are few sets of jointly sufficient variables, each of which is composed of many individually necessary variables (Mackie, 1974).

Of course, there are also many causal systems where this assumption does not hold; that is, where sufficiency is easier to compute than necessity. In such cases, people may shift to sufficiency judgments to estimate intervention effectiveness, and our model may be less applicable. But, even then, the necessity judgments described here may still function as a supplementary source of information about intervention effectiveness. Necessity-based causation judgments need not be the *only* method of approximating

intervention effectiveness; our claim is that token causation is one important cognitive tool that fills that role.

Descriptive boundary condition: cases of overdetermination and preemption The basic version of our model provides a tight fit to people’s judgments in conjunctive structures (Morris et al., 2019). In disjunctive structures with overdetermination or preemption, however, its predictions are inaccurate. These are cases in which two events are each sufficient to produce an outcome, and both happen (or one happens, and the other would have happened had the first failed). In these cases, neither event was necessary for the outcome; hence, if people are implementing the algorithm described here, they should consider neither event a cause. But in fact, people think they are causal (Gerstenberg et al., 2015; Halpern, 2016; Spellman & Kincannon, 2001). Moreover, in overdetermined cases, the abnormal inflation effect reverses—people judge rarer variables to be *less* causal (Gerstenberg & Icard, submitted; Icard et al., 2017). These findings are at odds with the basic version of our model.

Many accounts have been developed to explain people’s judgments in cases of overdetermination and preemption (Gerstenberg et al., 2015; Halpern, 2016; Halpern & Pearl, 2005; Icard et al., 2017; Lewis, 2000; Stephan & Waldmann, 2016). One possibility is that, even though our model accurately describes one important aspect of the functional design of causal judgments, causal judgment in overdetermined cases is explained by a different mechanism with a different functional purpose.

One explanation, for instance, comes from process theories. These assign causality to events that were connected to the outcome via some physical process (Dowe, 2000; Wolff, 2007; Wolff, Barbey, & Hausknecht, 2010). According to process theories, if two soldiers simultaneously shoot a prisoner in the head, both events were physically connected to the prisoner’s death—and hence both should rightly be considered a cause. Process theories thus offer a natural explanation for people’s intuitions in cases of overdetermination. The operation of process-based mechanism could be reconciled with our account if, as is often argued, people possess two different concepts of causation: one dependence concept and another process concept (Hall, 2004; Lombrozo, 2010). On

this view, our model would offer a functional account of people’s dependence concept, which is simply overshadowed by the process concept in cases of overdetermination/preemption. Despite their appeal, however, process theories have been criticized on various grounds (e.g., difficulty in dealing with cases of double prevention, Gerstenberg et al. 2014; Schaffer 2000, 2004).

Alternatively, it is possible that the basic function of the SAMPLE model is preserved in overdetermination cases, but with appropriate modifications to the simple algorithm described here. In other words, with some additional computation, people may be able to extend the sample-based method of estimating the effectiveness of interventions even from cases where no individual variable was necessary to produce an effect. We briefly sketch two possible approaches.

First, people might not only be able to assess whether any given variable setting was necessary to produce an effect; in some cases, they might invest additional effort to assess whether any given variable was sufficient, given certain background conditions (Icard et al., 2017; Woodward, 2006). As noted above, necessity judgments are an efficient way to approximate the effectiveness of interventions, but those approximations can be improved by also considering sufficiency. When two soldiers simultaneously shoot a prisoner in the head, both are sufficient for the prisoner’s death (holding constant certain background conditions). If people are making additional sufficiency judgments of this form, they might therefore judge the soldiers causal (Spellman & Kincannon, 2001). Potentially, an appropriate correction term applied to such judgments would allow them to accumulate information about general intervention effectiveness.

A related possibility is that people are not limited to judging whether a variable setting was necessary in the actual circumstances in which it occurred; in some cases, they may also judge whether it would have been necessary in slightly modified circumstances. For instance, if two soldiers simultaneously shoot a prisoner in the head, neither was necessary in the actual circumstances—but the first soldier would have been necessary had the second soldier not pulled the trigger (Halpern & Hitchcock, 2011;

Halpern & Pearl, 2005; Hitchcock, 2001; Lucas & Kemp, 2015).¹⁸ Exploiting this fact, people could (i) anchor on the actual state of affairs, (ii) adjust minimally to obtain a counterfactual state of affairs in which some variable setting would indeed have been necessary, and then (iii) apply our algorithm in the minimally counterfactual state of affairs. This would require the application of an additional correction term that adjusts for the relative frequency of the actual versus the counterfactual states of affairs. More concretely, by considering the counterfactual state in which the second soldier didn't shoot, people can sharpen their estimates of intervention effectiveness without having to wait to observe such an episode. Instead, they generate a simulated "observation" by adjusting minimally from the state of affairs they actually observed.¹⁹

Both of these approaches—the sufficiency analysis and the counterfactual necessity analysis—serve the same basic function as the SAMPLE model. By using a generative model of the situation to infer what would have happened in plausible altered circumstances, people may be generating additional "observations" of the causal system (Blanchard, Vasilyeva, & Lombrozo, 2017; Vasilyeva et al., 2016). Exploring these and other accounts of causal judgment in overdetermination and preemption remains an important area for further research.

8.2 Relationship to prior work on normality

Some prior treatments of token causation propose that it is closely related to the problem of finding good interventions. Hitchcock and Knobe (2009) suggested that the two concepts are linked, and argued that this linkage could explain the effects of normality on causal judgment (see also Woodward, 2003). This idea is at the core of our proposal. Our model elaborates and formalizes this idea, showing that a causation algorithm designed to identify good interventions would exhibit many observed effects of normality, as well as other empirical patterns.

¹⁸This explanation assumes that there are restrictions on which circumstances are permissible to modify. For details, see Halpern and Hitchcock (2011).

¹⁹This account would offer a functional explanation for the reversal of abnormal inflation in overdetermined cases (Icard et al., 2017). When generating the additional observation, people would presumably be influenced by the base rates of the antecedent events. Concretely, if the first soldier is more likely to shoot than the second, people should be more likely to generate an additional sample where the second soldier didn't shoot—and hence should consider the *more common* factor more causal.

Our model focuses exclusively on the effects of statistical abnormality on causal judgment—i.e. the fact that people are more likely to select rare events as causes in conjunctive settings. But the effects of normality go beyond statistics: People also believe that prescriptively abnormal events are more causal (Hitchcock & Knobe, 2009; Icard et al., 2017; Phillips et al., 2015). Agents who violate prescriptive norms, or objects which violate functional norms, are considered more of a cause of the resulting outcomes (Hitchcock & Knobe, 2009; Phillips & Kominsky, 2017). Prescriptive norms also apply to the supersession effect; an agent in a conjunctive setting is considered less causal if the *alternative* variable violated a prescriptive norm (Kominsky et al., 2015). Wherever statistical norms have an effect, prescriptive norms seem to have a similar effect (Bear & Knobe, 2017)

How can our model account for these effects? As noted above, our model is primarily a functional account of token causation; it is likely missing elements of the mechanism underlying real causal judgments. One salient element comes from Bear and Knobe (2017), who argue that, across many cognitive tasks, people replace estimates of a variable’s base rate—e.g., $Prob(C = 1)$, or $Prob(E = 1)$ —with an estimate of how “normal” the variable is, where “normal” integrates both statistical and prescriptive norms. If this is true, then prescriptive normality would have a similar effect as statistical normality in our model (Icard et al., 2017). Of course, it is also possible that there is some specific functional reason that causal judgments would incorporate prescriptive normality (Hitchcock & Knobe, 2009). Perhaps the relevant notion of “effective interventions” encompasses not only the probability that an intervention will be effective in general, but also a general estimate of its associated costs or benefits. We leave it to future work to investigate this and other related possibilities.

8.3 Future directions

Our model predicts the form that token causal judgments will take in a variety of causal structures. We find that it matches people’s judgments in a two-variable conjunctive structure; future work should investigate the quantitative form that people’s causal judgments take in other structures (e.g. structures with more variables,

or that combine conjunctive and disjunctive relationships). Future work should also test a key prediction of our model: People will use their causal judgments to help choose interventions. This could be tested by presenting people with a novel causal system where they can observe past episodes, make causal judgments, and then intervene to produce desired outcomes. Our model predicts that a person’s average judgment of the extent to which a variable C caused an outcome E will predict how often they choose C as an intervention to produce E .

Additionally, as noted above, our model admits of at least two interpretations. It could be interpreted as describing only the function of token causal judgments—occupying, in other words, Marr’s computational level of analysis (Marr, 1982). On this view, the computation being performed by token causation results in an approximation of intervention effectiveness, but the algorithm used to perform that computation may look different than what we have proposed. For instance, the basic functional design embodied by our model may be consistent with quite a different algorithm for generating estimates of intervention effectiveness, such as an algorithm that samples counterfactual scenarios (Gerstenberg et al., 2015; Hitchcock & Knobe, 2009; Icard et al., 2017; Kominsky et al., 2015; Phillips et al., 2015), or assesses the correlation between cause and effect across possible worlds (Quillien, 2019).

Alternatively, our model could be interpreted as a mechanistic description of how people make causal judgments: They first judge whether a variable was necessary, then apply the correction term, and so on (Marr’s algorithmic level of analysis). Refining our understanding of the particular algorithm that people employ during causal judgment is an important goal for future research.

9 Conclusion

Token causation and causal selection are central concepts in people’s everyday lives, but their form and function have remained mysterious. Here, we find that, if token causation is designed to approximate the effectiveness of potential interventions, it will take a form that matches observed empirical patterns. The nature of causal concepts may therefore be illuminated by focusing on their function (Woodward, 2014).

Acknowledgements

We thank Jonathan Kominsky and Max Kleiman-Weiner for their feedback on the manuscript. We also thank Jonathan Kominsky for sharing his data on supersession, and the Moral Psychology Research Lab for their advice and assistance. AM and FC were supported by Grant N00014-14-1-0800 from the Office of Naval Research to FC. TG was supported by the Center for Brains, Minds & Machines (CBMM), which is funded by the National Science Foundation's Science and Technology Center (Award CCF-1231216).

10 References

- Alicke, M. D. (2000). Culpable control and the psychology of blame. *Psychological bulletin*, 126(4), 556.
- Bear, A., & Knobe, J. (2017). Normality: Part descriptive, part prescriptive. *cognition*, 167, 25–37.
- Blanchard, T., Vasilyeva, N., & Lombrozo, T. (2017, aug). Stability, breadth and guidance. *Philosophical Studies*.
- Bramley, N. R., Dayan, P., Griffiths, T. L., & Lagnado, D. A. (2017). Formalizing neurath’s ship: Approximate algorithms for online causal learning. *Psychological review*, 124(3), 301.
- Bramley, N. R., Gerstenberg, T., Mayrhofer, R., & Lagnado, D. A. (in press). Time in causal structure learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*.
- Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, 104(2), 367–405.
- Cheng, P. W., & Novick, L. R. (1991). Causes versus enabling conditions. *Cognition*, 40, 83–120.
- Cheng, P. W., & Novick, L. R. (2005). Constraints and nonconstraints in causal learning: Reply to white (2005) and to luhmann and ahn (2005). *Psychological Review*, 112(3), 694–707.
- Chockler, H., Halpern, J. Y., & Kupferman, O. (2008). What causes a system to satisfy a specification? *ACM Transactions on Computational Logic*, 9(3), 20.
- Coenen, A., Rehder, B., & Gureckis, T. M. (2015). Strategies to intervene on causal systems are adaptively selected. *Cognitive psychology*, 79, 102–133.
- Collingwood, R. G. (2014). *An Essay on Metaphysics*. Martino Fine Books.
- Danks, D. (2017). Singular causation. In M. Waldmann (Ed.), *The oxford handbook of causal reasoning* (pp. 201–215). Oxford University Press.
- Dolan, R. J., & Dayan, P. (2013). Goals and habits in the brain. *Neuron*, 80(2), 312–325.

- Dowe, P. (2000). *Physical causation*. Cambridge, England: Cambridge University Press.
- Eiter, T., & Lukasiewicz, T. (2002). Complexity results for structure-based causality. *Artificial Intelligence*, 142(1), 53–89.
- Franklin, N. T., & Frank, M. J. (2018). Compositional clustering in task structure learning. *bioRxiv*, 196923.
- Gerstenberg, T., Goodman, N., Lagnado, D., & Tenenbaum, J. (2012). Noisy newtons: Unifying process and dependency accounts of causal attribution. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 34).
- Gerstenberg, T., Goodman, N., Lagnado, D., & Tenenbaum, J. (2014). From counterfactual simulation to causal judgment. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 36).
- Gerstenberg, T., Goodman, N. D., Lagnado, D. A., & Tenenbaum, J. B. (2015). How, whether, why: Causal judgments as counterfactual contrasts. In D. C. Noelle et al. (Eds.), *Proceedings of the 37th Annual Conference of the Cognitive Science Society* (pp. 782–787). Austin, TX: Cognitive Science Society.
- Gerstenberg, T., & Icard, T. F. (submitted). Expectations affect physical causation judgments.
- Gerstenberg, T., Peterson, M. F., Goodman, N. D., Lagnado, D. A., & Tenenbaum, J. B. (2017, oct). Eye-tracking causality. *Psychological Science*, 28(12), 1731–1744.
- Griffiths, T. L., & Tenenbaum, J. B. (2005). Structure and strength in causal induction. *Cognitive Psychology*, 51(4), 334–384.
- Griffiths, T. L., & Tenenbaum, J. B. (2009). Theory-based causal induction. *Psychological Review*, 116(4), 661–716.
- Hagmayer, Y., & Sloman, S. A. (2009). Decision makers conceive of their choices as interventions. *Journal of Experimental Psychology: General*, 138(1), 22–38.
- Hall, N. (2004). Two concepts of causation. In J. Collins, N. Hall, & L. A. Paul (Eds.), *Causation and counterfactuals*. MIT Press.

- Halpern, J. Y. (2016). *Actual causality*. MIT Press.
- Halpern, J. Y., & Hitchcock, C. (2011, June). Actual causation and the art of modeling. *arXiv:1106.2652 [cs]*. (arXiv: 1106.2652)
- Halpern, J. Y., & Hitchcock, C. (2015). Graded causation and defaults. *British Journal for the Philosophy of Science*, 66, 413–457.
- Halpern, J. Y., & Pearl, J. (2005). Causes and explanations: A structural-model approach. Part I: Causes. *The British Journal for the Philosophy of Science*, 56(4), 843–887.
- Hart, H. L. A., & Honoré, T. (1959/1985). *Causation in the law*. New York: Oxford University Press.
- Hesslow, G. (1988). The problem of causal selection. In D. J. Hilton (Ed.), *Contemporary science and natural explanation: Commonsense conceptions of causality* (pp. 11–32). Brighton, UK: Harvester Press.
- Hilton, D. J. (1990). Conversational processes and causal explanation. *Psychological Bulletin*, 107(1), 65–81.
- Hilton, D. J., & Slugoski, B. R. (1986). Knowledge-based causal attribution: The abnormal conditions focus model. *Psychological Review*, 93(1), 75–88.
- Hitchcock, C. (2001). The intransitivity of causation revealed in equations and graphs. *The Journal of Philosophy*, 98(6), 273–299.
- Hitchcock, C. (2009). Structural equations and causation: six counterexamples. *Philosophical Studies*, 144(3), 391–401.
- Hitchcock, C. (2012). Portable causal dependence: A tale of consilience. *Philosophy of Science*, 79(5), 942–951.
- Hitchcock, C. (2017). Actual causation: What’s the use? *Making a Difference: Essays on the Philosophy of Causation*, 5.
- Hitchcock, C., & Knobe, J. (2009). Cause and norm. *Journal of Philosophy*, 11, 587–612.
- Holyoak, K. J., Lee, H. S., & Lu, H. (2010). Analogical and category-based inference: A theoretical integration with bayesian causal models. *Journal of Experimental*

Psychology: General, 139(4), 702.

Icard, T. F., Kominsky, J. F., & Knobe, J. (2017). Normality and actual causal strength. *Cognition*, 161, 80–93.

Kahneman, D., & Miller, D. T. (1986). Norm theory: Comparing reality to its alternatives. *Psychological Review*, 93(2), 136–153.

Kahneman, D., & Tversky, A. (1982). The simulation heuristic. In D. Kahneman & A. Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases* (pp. 201–208). New York: Cambridge University Press.

Kominsky, J. F., Phillips, J., Gerstenberg, T., Lagnado, D. A., & Knobe, J. (2015). Causal superseding. *Cognition*, 137, 196–209.

Lagnado, D. A., Gerstenberg, T., & Zultan, R. (2013). Causal responsibility and counterfactuals. *Cognitive science*, 37(6), 1036–1073.

Lewis, D. (1973). Causation. *The Journal of Philosophy*, 70(17), 556–567.

Lewis, D. (1979). Counterfactual dependence and time’s arrow. *Noûs*, 13(4), 455–476.

Lewis, D. (2000). Causation as influence. *The Journal of Philosophy*, 97(4), 182–197.

Lombrozo, T. (2006, October). The structure and function of explanations. *Trends in Cognitive Sciences*, 10(10), 464–470.

Lombrozo, T. (2010). Causal-explanatory pluralism: How intentions, functions, and mechanisms influence causal ascriptions. *Cognitive Psychology*, 61(4), 303–332.

Lu, H., Yuille, A. L., Liljeholm, M., Cheng, P. W., & Holyoak, K. J. (2008). Bayesian generic priors for causal learning. *Psychological Review*, 115(4), 955–984.

Lucas, C. G., & Kemp, C. (2015). An improved probabilistic account of counterfactual reasoning. *Psychological Review*, 122(4), 700–734.

Luhmann, C. C., & Ahn, W.-k. (2005, July). The Meaning and Computation of Causal Power: Comment on and. *Psychological review*, 112(3), 685–707.

Mackie, J. L. (1974). *The cement of the universe*. Oxford: Clarendon Press.

Marr, D. (1982). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. New York, NY, USA: Henry Holt and Co., Inc.

- McGill, A. L., & Tenbrunsel, A. E. (2000). Mutability and propensity in causal selection. *Journal of Personality and Social Psychology*, 79(5), 677-689.
- Meder, B., Gerstenberg, T., Hagmayer, Y., & Waldmann, M. R. (2010). Observing and intervening: Rational and heuristic models of causal decision making. *Open Psychology Journal*, 3, 119–135.
- Morris, A., Phillips, J. S., Gerstenberg, T., & Cushman, F. (2019). Quantitative causal selection patterns in token causation. <https://psyarxiv.com/upv8t>.
- Pearl, J. (1999). Probabilities of causation: three counterfactual interpretations and their identification. *Synthese*, 121(1-2), 93–149.
- Pearl, J. (2000). *Causality: Models, reasoning and inference*. Cambridge, England: Cambridge University Press.
- Phillips, J., & Cushman, F. (2017). Morality constrains the default representation of what is possible. *Proceedings of the National Academy of Sciences*, 114(18), 4649–4654.
- Phillips, J., & Kominsky, J. (2017). Causation and norms of proper functioning: Counterfactuals are (still) relevant. In G. Gunzelmann, A. Howes, T. Tenbrink, & E. Davelaar (Eds.), *Proceedings of the 39th Annual Conference of the Cognitive Science Society* (pp. 931–936). Austin, TX: Cognitive Science Society.
- Phillips, J., Luguri, J. B., & Knobe, J. (2015). Unifying morality's influence on non-moral judgments: The relevance of alternative possibilities. *Cognition*, 145, 30–42.
- Quillien, T. (2019). When do we think that x caused y?
- Rottman, B. M., & Hastie, R. (2013). Reasoning about causal relationships: Inferences on causal networks. *Psychological Bulletin*.
- Samland, J., & Waldmann, M. R. (2016). How prescriptive norms influence causal inferences. *Cognition*, 156, 164–176.
- Sartorio, C. (2012, Dec). Two wrongs do not make a right: Responsibility and overdetermination. *Legal Theory*, 18(04), 473-490.
- Schaffer, J. (2000). Causation by disconnection. *Philosophy of Science*, 67(2), 285–300.

- Schaffer, J. (2004). Causes need not be physically connected to their effects: The case for negative causation. In C. R. Hitchcock (Ed.), *Contemporary debates in philosophy of science* (pp. 197–216). Blackwell.
- Schultz, W. (2002). Getting formal with dopamine and reward. *Neuron*, 36(2), 241–263.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., . . . others (2016). Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587), 484–489.
- Sloman, S. A. (2005). *Causal models: How people think about the world and its alternatives*. Oxford University Press, USA.
- Sloman, S. A., Barbey, A. K., & Hotaling, J. M. (2009). A causal model theory of the meaning of cause, enable, and prevent. *Cognitive Science*, 33(1), 21–50.
- Sloman, S. A., & Hagmayer, Y. (2006). The causal psychology of choice. *Trends in Cognitive Sciences*, 10(9), 407–412.
- Sloman, S. A., & Lagnado, D. A. (2005). Do we ‘do’? *Cognitive Science*, 29(1), 5–39.
- Spellman, B. A. (1997). Crediting causality. *Journal of Experimental Psychology: General*, 126(4), 323–348.
- Spellman, B. A., & Kincannon, A. (2001). The relation between counterfactual ("but for") and causal reasoning: Experimental findings and implications for jurors' decisions. *Law and Contemporary Problems*, 64(4), 241–264.
- Stephan, S., & Waldmann, M. R. (2016). Answering causal queries about singular cases. In *Proceedings of the 38th annual conference of the cognitive science society* (pp. 2795–2801).
- Stephan, S., & Waldmann, M. R. (2018). Preemption in singular causation judgments: A computational model. *Topics in cognitive science*, 10(1), 242–257.
- Steyvers, M., Tenenbaum, J. B., Wagenmakers, E.-J., & Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science*, 27(3), 453–489.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction* (Vol. 1)

(No. 1). MIT press Cambridge.

Tian, J., & Pearl, J. (2000). Probabilities of causation: Bounds and identification.

Annals of Mathematics and Artificial Intelligence, 28(1-4), 287–313.

Vasilyeva, N., Blanchard, T., & Lombrozo, T. (2016). Stable causal relationships are

better causal relationships. In A. Papafragou, D. Grodner, D. Mirman, &

J. C. Trueswell (Eds.), *Proceedings of the 38th Annual Conference of the Cognitive Science Society* (pp. 2663–2668). Austin, TX: Cognitive Science Society.

Walsh, C. R., & Sloman, S. A. (2011). The meaning of cause and prevent: The role of causal mechanism. *Mind & Language*, 26(1), 21–52.

Wells, G. L., & Gavanski, I. (1989). Mental simulation of causality. *Journal of personality and social psychology*, 56(2), 161.

Willemsen, P., & Kirfel, L. (2018). Recent empirical work on the relationship between causal judgements and norms. *Philosophy Compass*, e12562.

Wolff, P. (2007). Representing causation. *Journal of Experimental Psychology: General*, 136(1), 82–111.

Wolff, P., Barbey, A. K., & Hausknecht, M. (2010). For want of a nail: How absences cause events. *Journal of Experimental Psychology: General*, 139(2), 191–221.

Woodward, J. (2003). *Making things happen: A theory of causal explanation*. Oxford, England: Oxford University Press.

Woodward, J. (2006). Sensitive and insensitive causation. *The Philosophical Review*, 115(1), 1–50.

Woodward, J. (2011). Psychological studies of causal and counterfactual reasoning. In C. Hoerl, T. McCormack, & S. R. Beck (Eds.), *Understanding counterfactuals, understanding causation: Issues in philosophy and psychology*. Oxford: Oxford University Press.

Woodward, J. (2014, December). A Functional Account of Causation; or, A Defense of the Legitimacy of Causal Thinking by Reference to the Only Standard That Matters—Usefulness (as Opposed to Metaphysics or Agreement with Intuitive Judgment). *Philosophy of Science*, 81(5), 691–713.

11 Appendix

i: Introduction to structural equation modeling

A structural equation model is a formal description of a causal system that allows a decision-maker to compute the downstream effects of interventions (Pearl, 2000). It is composed of three parts: a set of exogenous variables, a set of endogenous variables, and functions that relate them to each other. The exogenous variables represent all the elements of the causal system whose antecedents the decision-maker does not want to explicitly model. For instance, one might create a structural equation model of forest fires (Figure 6A) where whether or not an arsonist is present is an exogenous variable: the decision-maker does not model the variables that determine that. Instead, for convenience, she simply assigns a prior probability to it (e.g. there might be a 10% chance that, in any observation of the system, an arsonist is present).

In contrast, the endogenous variables represent the elements that the decision-maker *does* want to model. The values of the endogenous variables are determined by the other variables in the causal system. For instance, in our toy example, there is a forest fire if and only if there is oxygen, dry ground, and a lit match. This relationship can be captured by describing the variable “forest fire” (FF) as a function of the variables “oxygen” (O), “dryness” (D), and “lit match” (LM). In this case, $FF = \min(O, D, LM)$; O, D, and LM must all equal 1 for FF to equal 1. Of course, the decision-maker does not typically know the values of all the exogenous variables. Hence, to compute the probability that an *endogenous* variable takes on a value (e.g. $Prob(FF = 1)$), she simply identifies the settings of the *exogenous* variables that would produce such a value (e.g. $O = 1, D = 1, A = 1$), and computes the probability of those settings. This technique allows her to simultaneously represent the deterministic forces that produce an event, and also the uncertainty surrounding its actual occurrence (i.e., a probability distribution).

Structural equation models support two operations that are essential for defining the effectiveness of an intervention: intervening and conditioning. An intervention selectively alters the value of a single variable, severing its ties to antecedent variables

Figure 6. (A) An example structural equation model of forest fires. Conceptually, $A = 1$ indicates that there is an arsonist; $D = 1$ indicates that the ground is dry; $LM = 1$ indicates that someone lit a match; $O = 1$ indicates that there is oxygen; and $FF = 1$ indicates that there is a forest fire. A , D , and O are exogenous; LM and FF are endogenous. In this toy model, we assume that there is a lit match if and only if there is an arsonist, and there is a forest fire if and only if someone lights a match, the ground is dry, and there's oxygen in the air. (Conventionally, exogenous variables are not depicted in graphical representations – but we show them here for exposition.) (B) Procedure for computing the probability of fire after lighting a match, conditioned on FF currently being zero. First, you update the probabilities of the exogenous variables to account for $FF = 0$. Second, you intervene to fix $LM = 1$. Third, you sum the conditional probabilities of the exogenous variable settings that would, post-intervention, produce $FF = 1$.

that may have influenced it. Structural equation models formalize this with the notion of a *mutilated model*. Intervening to set a variable V to value v in structural equation model M , the decision-maker considers a derivative model $M_{V=v}$ in which the function that determines V is replaced with $V = v$. (This process is often referred to as applying the “do” operator; see Pearl, 2000; Sloman & Lagnado, 2005.) Then, to determine the effects of this hypothetical intervention, the intervener computes the probabilities that other variables would take certain values in this mutilated model. For instance, to compute the probability that there would be a fire if you lit a match, we compute the probability that $FF = 1$ in the mutilated model with LM fixed to 1 (notated $Prob(FF_{LM=1} = 1)$).

The second operation, conditioning, involves updating the probabilities in the

causal model based on knowledge of variable values. For instance, one could compute the probability that there would be a forest fire, conditioned on your knowledge that the ground is dry ($Prob(\mathbf{FF} \mid \mathbf{D}))$).

Combining these operations, we can ask a question critical for our purposes: What is the probability of a forest fire given that you lit a match, conditioned on the fact that there is currently no forest fire? Formally, we notate this $Prob(\mathbf{FF}_{\mathbf{LM}=1} \mid \mathbf{FF} = 0)$. This operation proceeds in three steps (Pearl, 2000; Figure 6B). First, the decision-maker computes the conditional probabilities of the exogenous variables given her knowledge of the system’s current state (i.e. $Prob(\mathbf{A} = 1 \mid \mathbf{FF} = 0)$, $Prob(\mathbf{D} = 1 \mid \mathbf{FF} = 0)$, etc.). Second, she mutilates the model to simulate the intervention (fixing $\mathbf{LM} = 1$). Finally, she computes the probability of the outcome ($\mathbf{FF} = 1$) in the mutilated model, using the conditional probabilities from step one. In this way, structural equation models allow her to compute the effects of interventions while incorporating her knowledge of specific situations.²⁰

ii: Why our measure of intervention effectiveness conditions on $\mathbf{E} = 0$

Our definition of the general effectiveness of an intervention — $GE(C \rightarrow E)$ — conditions on the absence of the outcome; it is defined as $Prob(E_{C=1} = 1 \mid E = 0)$. In other words, the usefulness of an intervention to bring about an outcome is measured by situations where the outcome is currently absent. The purpose of this section is to elaborate on why we defined intervention effectiveness this way. To reiterate, the argument is that, in decision-making situations where a person is trying to bring about a goal ($E = 1$), the one thing that the person almost always knows is that the goal does not currently obtain (i.e. before the intervention, $E = 0$). If she knows that the goal already obtains, then she would not be trying to bring it about, and so would not consult her representation of the general effectiveness of potential interventions. Since the representation is only used in situations where $E = 0$, it is beneficial to condition the representation on that knowledge. Otherwise, distortions can occur.

For instance, one distortion is that a perfectly effective intervention can be

---

²⁰Pearl (2000) calls this process “counterfactual inference”. For a more formal treatment of it, and of structural equation modeling in general, see Pearl (2000).

represented as ineffective. Imagine a lightbulb that turns on if and only if two switches are in the same orientation (up or down). Intuitively, if you want to turn the light on, flipping each switch is a perfectly effective intervention. This is because, in situations where you want to turn the light on, the light is always currently off – and so flipping either switch is guaranteed to turn the light on. A measure of general effectiveness that conditions on the absence of the desired outcome will capture this intuition;

$$Prob(\text{Light}_{\text{Flip}} = \text{on} \mid \text{Light} = \text{off}) = 1.$$

But if you did not condition on the light currently being off, then you would consider each switch to be a less effective intervention, because you would be incorporating cases where the lightbulb is currently on (and hence the switch turns the light off). Imagine that our model of general effectiveness was $Prob(\text{Light}_{\text{Flip}} = \text{on}) - Prob(\text{Light} = \text{on})$ — a salient alternative measure. In cases when the light was usually on, we would consider flipping the switches to be ineffective. This is the distortion caused by not conditioning on the absence of the outcome.

There are, of course, disadvantages of conditioning on $E = 0$. Our measure is useful for *selecting* interventions in cases where $E = 0$. But, because it does not consider the cases in which E is already present, it is unhelpful in deciding *whether* to intervene. If an outcome is likely to occur naturally, and there are costs to intervening, then $GE(C \rightarrow E)$ will recommend choosing interventions when the optimal decision is actually to do nothing. In contrast, a measure like $Prob(E_{C=1} = 1) - Prob(E = 1)$ is more suited to decide whether to intervene, because it incorporates the possibility that the outcome will happen without intervention. Our measure is certainly not the only quantity needed in planning; but it would be a useful quantity for causation judgments to estimate.²¹

One final point: You might wonder, why don't we condition on the absence of the potential intervention also? In other words, why not condition on $C = 0$? The reason is

---

²¹Interestingly, our measure of $GE(C \rightarrow E)$, $Prob(E_{C=1} = 1 \mid E = 0)$, is closely related to the alternate measure $Prob(E_{C=1} = 1) - Prob(E = 1)$. If interventions cannot decrease the probability of E (the “monotonicity” assumption (Pearl, 1999)), then $Prob(E_{C=1} = 1) - Prob(E = 1) = GE(C \rightarrow E)Prob(E = 0)$. Hence, a quantity useful for deciding whether to intervene can be easily derived from the quantity that our algorithm estimates.

that, across all decision situations, it is not generally true that $C = 0$. For instance, consider choosing interventions to help a student get into college. We should not assume that the student has not already improved their writing, or joined the debate team, or pulled any of the scholastic levers that could serve as potential interventions. Decision situations will vary wildly in which potential interventions have already been employed, and hence the representation of the general, situation-agnostic effectiveness of an intervention should not condition on the intervention being absent.

iii: On the assumption of binary variables

For simplicity, the derivation of our model assumed that all variables in a decision-maker's model—causes and effects—were binary. Here, we consider the restrictiveness of this assumption. For our derivation to hold, it is not required that variables actually be binary; all that is required is that the decision-maker *treat* the variables as binary in the relevant decision context. For instance, consider a categorical outcome with more than two categories, like a letter grade on a test. The variable **grade** has five possible values A, B, C, D, F , and yet it can still be treated as binary if the decision-maker's goal is to get an A on the test. In that case, the relevant outcomes are **grade** = A and **grade** $\neq A$. (It's not necessary for the goal to be a single outcome setting; if the decision-maker's goal was to get an A or B, the relevant outcomes would be **grade** $\in \{A, B\}$ and **grade** $\notin \{A, B\}$.) Similarly, a continuous variable — say, the test's numerical grade — can also be treated as binary if the decision-maker's goal is to get above (or below) a threshold. In that case, the relevant outcomes are **grade** $> threshold$ and **grade** $\leq threshold$. (The same considerations apply to categorical or continuous causes; so long as the decision-maker conceptualizes potential interventions as binary, our analysis remains unchanged.)

Difficulties only arise when decision-makers genuinely conceptualize variables as continuous in the relevant decision context. For instance, a decision-maker's goal may be to maximize the numerical grade they receive on a test. In that case, our definition of an effective intervention is inadequate. Instead of maximizing the probability of the outcome's occurrence (i.e. $Prob(E_{C=1} = 1 \mid E = 0)$), a decision-maker must maximize

the expected value of the outcome ($EV[E_{C=1}]$), perhaps conditional on the outcome currently occupying an undesired state. Although possible that the principles described in this paper will apply to such cases, this question is currently unresolved.

iv: Proof of Proposition 1

In this section, we prove the relationship stated in Proposition 1: that the general effectiveness of intervening on C to produce E can be estimated by determining necessity in cases where E is present, and applying a correction term. Here, instead of using $GE(C \rightarrow E)$ to denote the general effectiveness, we will use its more common name: the probability of enablement, denoted $PE(C, E)$. To restate it formally:

$$PE(C, E) \approx \frac{1}{k} \sum_{i=1}^k TC_{SAMPLE}(C \rightarrow E, u_i)$$

where:

$$TC_{SAMPLE}(C \rightarrow E, u_i) = \begin{cases} \frac{Prob(C=0) Prob(E=1)}{Prob(C=1) Prob(E=0)} & \text{if } C \text{ was necessary for } E \text{ in } u_i \\ 0 & \text{otherwise} \end{cases}$$

The proof draws heavily on Pearl (1999). Consider a causal system with at least two variables, C and E . We have already defined the probability of enablement for C on E as $PE(C, E) = Prob(E_{C=1} = 1 \mid E = 0)$. It will be useful to define one other value: the probability of disablement for C on E , or $PD(C, E)$. This is, intuitively, the opposite of the probability of enablement. It is the probability that, given that the outcome E is present, turning C off would turn E off. Formally:

$$PD(C, E) = Prob(E_{C=0} = 0 \mid E = 1) \tag{1}$$

Intuitively, the probability of disablement can be estimated by assessing the necessity of C in situations where $E = 1$. We state this formally.

Proposition 2. Suppose $u_1, \dots, u_k$ are randomly sampled observations from the joint

probability distribution over all exogenous variables, given $E = 1$. As k increases, the following holds:

$$PD(C, E) \approx \frac{1}{k} \sum_{i=1}^k I(C \text{ was necessary for } E \text{ in } u_i)$$

where I is a function that equals 1 if its argument is true, and 0 otherwise. (The proof follows immediately from the law of large numbers.)

This is, of course, precisely the sampling procedure that we want to use. Our task, then, is to relate the probability of enablement (our target) to the probability of disablement (which can be estimated via this procedure). This is accomplished by one crucial assumption. PE and PD can be easily transformed into each other, so long as the effect on E from intervening on C is not dependent on the current value of C . In other words, $Prob(E_{C=1} = 1, E_{C=0} = 0 \mid C = c)$ must equal $Prob(E_{C=1} = 1, E_{C=0} = 0)$. This is the “no-confounding”, or “exogeneity” assumption (Pearl, 2000), which we discussed above.

If the no-confounding assumption holds, then the probabilities of enablement and disablement are related in the following way (Pearl, 1999).

Proposition 3. Suppose that

$Prob(E_{C=1} = 1, E_{C=0} = 0 \mid C = c) = Prob(E_{C=1} = 1, E_{C=0} = 0)$. Then:

$$PE(C, E) = \frac{P(C = 0)P(E = 1)}{P(C = 1)P(E = 0)} PD(C, E)$$

Proof. By the law of total probability,

$$\begin{aligned} PE(C, E) &= Prob(E_{C=1} = 1 \mid E = 0) = \\ &= Prob(E_{C=1} = 1 \mid E = 0, C = 0) * Prob(C = 0 \mid E = 0) \\ &\quad + Prob(E_{C=1} = 1 \mid E = 0, C = 1) * Prob(C = 1 \mid E = 0) \end{aligned}$$

We will make use of the fact that $Prob(E_{C=c} = e \mid E \neq e, C = c) = 0$. In words, if C is already equal to a value c , then intervening to set it to c cannot change the value

of E (Pearl, 2000). (Another way to write this, used below, is that

$Prob(E = e, C = c) = Prob(E_{C=c} = e, C = c)$.) Hence:

$$\begin{aligned}
 PE(C, E) &= \\
 Prob(E_{C=1} = 1 \mid E = 0, C = 0) * Prob(C = 0 \mid E = 0) &= \\
 Prob(E_{C=1} = 1, E = 0, C = 0) * \frac{Prob(C = 0 \mid E = 0)}{Prob(E = 0, C = 0)} &= \\
 Prob(E_{C=1} = 1, E = 0, C = 0) * \frac{1}{Prob(E = 0)} &= \\
 Prob(E_{C=1} = 1, E_{C=0} = 0, C = 0) * \frac{1}{Prob(E = 0)} &= \\
 Prob(E_{C=1} = 1, E_{C=0} = 0 \mid C = 0) * \frac{Prob(C = 0)}{Prob(E = 0)} &=
 \end{aligned}$$

By identical logic:

$$PD(C, E) = Prob(E_{C=1} = 1, E_{C=0} = 0 \mid C = 1) * \frac{Prob(C = 1)}{Prob(E = 1)}$$

By the no-confounding assumption,

$$Prob(E_{C=1} = 1, E_{C=0} = 0 \mid C = 0) = Prob(E_{C=1} = 1, E_{C=0} = 0 \mid C = 1)$$

Hence,

$$PE(C, E) * \frac{Prob(E = 0)}{Prob(C = 0)} = PD(C, E) * \frac{Prob(E = 1)}{Prob(C = 1)}$$

and

$$PE(C, E) = \frac{P(C = 0)P(E = 1)}{P(C = 1)P(E = 0)} PD(C, E)$$

□

Combining Propositions 2 and 3 completes the proof of Proposition 1.

v: Arguments for normalization

There are two reasons people might normalize their responses. One reason is pragmatic. The output of TC_{SAMPLE} has an unbounded upper range, but we ask people to rate causality on a bounded Likert scale. Moreover, in conversations involving

two salient contributing variables, people may assume that they are being asked to compare these variables to each other. It is thus possible that, although people employ the unnormalized output when privately interacting with the causal system, they often report normalized results when asked to explicitly judge a variable’s causal status in a social setting.

But normalization may also serve a deeper purpose. As described above, a key function of estimating the general effectiveness of potential interventions would be to narrow down the set of interventions to consider at decision time. In other words, causality-driven estimates of $GE(C \rightarrow E)$ may help a decision-maker construct a “choice set” of plausible interventions available at decision time (Phillips & Cushman, 2017). Normalizing causal judgments would support this function. By comparing potential interventions to their alternatives, normalization can exaggerate the differences between them and highlight the top candidates to consider at decision time.