

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

How People Make Causal Judgments about Unprecedented Societal Events

Permalink

<https://escholarship.org/uc/item/2d84w5fp>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 43(43)

ISSN

1069-7977

Authors

Amemiya, Jamie
Heyman, Gail D.
Walker, Caren M.

Publication Date

2021

Peer reviewed

How People Make Causal Judgments about Unprecedented Societal Events

Jamie Amemiya, Gail D. Heyman, & Caren M. Walker

(jamemiya@ucsd.edu; gheyman@ucsd.edu; carenwalker@ucsd.edu)

Department of Psychology, University of California, San Diego

9500 Gilman Drive

La Jolla, CA 92093 USA

Abstract

Counterfactual theories of causal judgment propose that people infer causality between events by comparing an actual outcome with what would have happened in a relevant alternative situation. If the candidate cause is “difference-making”, people infer causality. This framework has not been applied to people’s judgments about unprecedented societal events (e.g., global pandemics), in which people have limited causal knowledge (e.g., about effective policies). In these contexts, it is less clear how people reason counterfactually. This study examined this issue. Participants judged whether a mandatory evacuation reduced population bite rates during a novel insect infestation. People tended to rely on prior causal knowledge, unless data from *close alternatives* (i.e., structurally similar counterfactuals) provided counterevidence. There were also notable individual differences, such that some people privileged prior knowledge regardless of the available counterevidence or privileged *far alternatives* (i.e., structurally distinct counterfactuals), which may have implications for understanding public disagreement about policy issues.

Keywords: causal judgments; counterfactual reasoning; public policy

Introduction

Publicly available data indicates that the United States had the highest number of coronavirus cases in the world in 2020. While this *outcome* is evident, the *causes* of this outcome remain a frequent topic of public debate. Was this outcome affected by the delay in a mandatory lockdown? Did the structure of the United States’ health care system play a role?

According to counterfactual theories of causal judgment, people engage in counterfactual simulations to answer these types of causal questions (Gerstenberg et al., 2021; Kominsky & Phillips, 2019; Lewis, 1973; Mackie, 1974; Woodward, 2005). Specifically, people compare the actual, known outcome to a relevant alternative situation that informs what would have happened had things been different. If this counterfactual simulation reveals that the candidate cause is “difference-making”, people endorse it as causal.

A key assumption of contemporary counterfactual theory is that “people already have access to a generative model of the domain” (Gerstenberg et al., 2021). In other words, it is assumed that people have the requisite prior knowledge to generate appropriate counterfactuals. This work has, in turn, focused on domains in which people have robust prior causal knowledge. For instance, many studies have examined physical causation, in which participants can rely on intuitive physics to simulate counterfactual situations. To illustrate, a recent eye-tracking study found that, when determining

whether Ball A caused Ball B to go through a gate, people visualize the path that Ball B *would have* taken had it not been hit by Ball A (Gerstenberg et al., 2017). Other work has applied a counterfactual framework to understand people’s causal reasoning about common life experiences, including studying for an exam, taking pain medication, and preparing food (Byrne, 2016).

It is less clear how people make counterfactual and causal judgments about *unprecedented* events that have never happened before, for example, about whether a particular public policy was effective in controlling the spread of a new pandemic. Indeed, while people have a lifetime of experience observing and manipulating objects under various conditions, they lack such rich experience with pandemics. Despite this, it is notable that in almost any news outlet, there is plenty of speculation about which policies are responsible for a country’s outcomes (Beauchamp, 2020; Berlinger, 2020). Moreover, prior research finds that both lay individuals and historians frequently consider how things might have been different when making causal and moral judgments about historical events (Markman et al., 2008; Nolan, 2013; Tetlock & Belkin, 1996). Thus, we do not consider *whether* people make counterfactual and causal judgments about unprecedented events, but rather *how* they do so. In the present study, we examine how a counterfactual framework of causal judgment may apply to these scenarios. Below, we propose three possible accounts:

Account 1: Relying on (Limited) Prior Knowledge

One possibility is that people simulate counterfactuals in the same way they do when reasoning about familiar events by drawing on the limited knowledge that they do have. For example, when judging whether a delay in initiating a lockdown increased coronavirus cases, people may reason that staying away from others generally prevents the spread of disease and, to the extent that people follow the order, an earlier lockdown would have prevented infections. In this way, people would not seek data beyond the outcome of interest (e.g., infections rates in the United States) to make a causal judgment, as they can construct the relevant counterfactual world on their own. This tendency to rely on prior knowledge also aligns with alternative mechanistic accounts of causal judgment, which propose that people seek evidence for an underlying mechanism (e.g., transfer of causal force) when inferring causality (e.g., Wolff & Thorstad, 2017).

Account 2: Relying on Close Alternatives

A second possibility is that people may privilege data from *close alternatives*, or structurally similar cases (e.g., other societies with similar demographics), to generate counterfactuals. Indeed, this appears to be a common practice in journalism on the pandemic. For example, when speculating about the causes of a target country's outcomes, news outlets typically report *other* countries' coronavirus case trajectories in the same graph, and cases are often reported in units that allow for between-country comparison (e.g., reporting number of cases per million residents; Connolly, 2020). In line with this *close alternatives* account, countries that are graphed together often have a similar societal structure, including demographic and economic characteristics (e.g., comparing western European countries to one another). Thus, it may be that people rely on close alternatives when making causal judgments about unprecedented events, in that they rely on data from other societies that are structurally similar, but differ in the candidate cause of interest (e.g., policies).

Account 3: Individual Variability

A third possibility is that there are substantial individual differences, such that some people privilege their own causal knowledge and others privilege close alternatives. Further individual variability may come from differences in *which* alternatives people perceive to be relevant comparison cases. Indeed, the fact that there is widespread disagreement in many countries about whether their current administration handled the pandemic well (Mordecai & Connaughton, 2020) is suggestive that there may be notable individual differences in people's counterfactual reasoning and subsequent causal judgments.

The Present Study

The present study aimed to answer which of these three accounts best describes how people tend to make causal judgments about unprecedented societal events. Rather than manipulating participants' judgments about the pandemic, which may be difficult due to people's strong prior beliefs, participants were asked to reason about an unprecedented event that was conceptually similar.

Specifically, participants read a news article about a novel insect infestation in the target town, Hillsbrook, Pennsylvania, and were asked whether a delay in mandatory evacuation caused an increase in population bite rates and to explain why they made that judgment. Prior to these causal judgments, participants in the *Baseline* condition were only presented data on population bite rates from Hillsbrook. We purposely designed the bite rate trajectory to present ambiguous evidence for whether or not the evacuation was causal; specifically, there was a slowing down of bite rates once the evacuation was ordered on day 5 (perhaps suggesting a causal effect) but bite rates still showed a slight upward trend (perhaps suggesting no effect).

Participants in the other two conditions (*Causal* and *Not Causal*) were presented this baseline data, and were also presented with data from two other Pennsylvania towns (i.e., close alternatives) and two towns from developing Asian countries (i.e., far alternatives). Critically, participants were informed that these other four towns had *immediate* evacuations, and thus could be used to generate counterfactuals about what would have happened had Hillsbrook not had the delay. In the *Causal* condition, the other Pennsylvania towns had low bite rates and thus supported a causal effect of the delayed evacuation. In the *Not Causal* condition, the other Pennsylvania towns had essentially the same trajectory of bite rates as Hillsbrook, suggesting that the delayed evacuation was not difference-making. See Figure 1 for the data presented in each of the three conditions. To test whether people privilege these close alternatives, data from the Asian towns supported the *opposite* causal judgment for each condition.

To determine which of the three accounts best describes people's reasoning, we leveraged participants' causal judgments *and* their explanations. If participants were *only* relying on their limited prior knowledge (Account 1), then their judgments should be insensitive to the type of data that they are presented, and they should only mention causal mechanism information in their explanations (i.e., their beliefs about the effectiveness of evacuations). We assumed that participants would generally believe that evacuations are effective for avoiding poisonous insects, and thus would endorse that the delayed evacuation caused high bite rates if they were relying exclusively on their prior knowledge. Alternatively, if participants privilege close alternatives (Account 2), their causal judgments should vary as a function of the data that they are provided, and they should specifically reference the Pennsylvania towns—but *not* the Asian towns—in their explanations. Finally, Account 3 predicts that people's causal judgments and explanations should be linked in a coherent manner, but that there will be individual differences in whether prior causal knowledge, close alternatives, or far alternatives are privileged.

Method

Participants

Participants were 64 undergraduate students (64% female, 31% male; 44% Asian, 19% White, 16% mixed race, 11% Latinx, 8% Middle Eastern, 2% Black) who attended a large, public university in the West Coast region of the United States.

Procedure

Participants completed the study online via Qualtrics and were assigned to one of three, between-subject conditions: (1) *Baseline* (no data beyond the target town), (2) *Causal* (data from close alternatives—i.e., other Pennsylvania towns—suggesting that the candidate cause is causal), or (3) *Not Causal* (data from close alternatives suggesting the candidate cause is *not* causal).

All conditions began with the same cover story regarding a recent infestation of a poisonous novel insect:


In Hillsbrook, Pennsylvania, there was an infestation of a novel flying insect, Zorphax seymora. Scientists found out right away that the insect’s bite was highly poisonous to humans. Scientists told Hillsbrook to order an immediate evacuation to shelters further away from the infestation.

All participants were then told about a delay in intervention, a mandatory evacuation:


Due to unexpected technical issues, Hillsbrook got the evacuation message five days after it was sent. Once the message was received, Hillsbrook ordered the evacuation.

Participants in the *Baseline* condition were then shown the population bite rates of the target town (see Figure 1).

baseline



causal



not causal




Figure 1: Study conditions.

In the *Causal* and *Not Causal* conditions, the story described four additional towns (two towns that were close in location, i.e., also in Pennsylvania, and two towns from developing countries in another continent, i.e., Asia) that had the same insect infestation but had an *immediate* evacuation:

The same insect infestation also happened in the following towns: Allensville, Pennsylvania, Johnston, Pennsylvania, Kapung, Indonesia, and Dima, Malaysia. In all of these other towns, they ordered evacuations to shelters right away with no delay.

In the *Causal* condition, participants were presented data in which the close alternatives (Allensville, Pennsylvania and Johnston, Pennsylvania) suggested that the delayed evacuation made a difference, while the far alternatives (Dima, Malaysia and Kapung, Indonesia), suggested that delayed evacuation did *not* make a difference. In the *Not Causal* condition, the same exact data were presented but the labels were reversed, thus supporting the opposite inference.

Participants were then asked to make a *causal judgment*, specifically, “Did the delayed evacuation cause the high rates of Hillsbrook residents being bitten?” They responded on a scale from 0 to 100, with the anchors being “definitely no” at 0, “maybe” at 50, and “definitely yes” at 100. Following this, they were asked to explain their answer (“Why do you think that? Please write 1-2 sentences.”).

Participants also made a *counterfactual judgment*: “Imagine instead there was an immediate evacuation in Hillsbrook. Would it have led to a drop of at least 5% of Hillsbrook residents being bitten?” Participants rated their answers on the same response scale and were also asked to provide open-ended justifications for their responses. The order of presentation of the causal and counterfactual questions was counterbalanced across participants.

Results

Causal judgments

Participants’ causal judgments by condition are graphed in Figure 2, in which higher scores indicate stronger agreement that the delayed evacuation caused the high bite rates in Hillsbrook. A one-way ANOVA indicated that there was an overall effect of condition, $F(2) = 7.11, p = .002, \eta^2 = .19$. Post hoc pairwise comparisons indicated that participants in the *Baseline* condition were more likely to endorse the delayed evacuation as causal than participants in the *Not Causal* condition, $t(40) = 3.55, p < .001$, but no more likely than participants in the *Causal* condition, $t(39) = 1.07, p = .29$. Critically, participants in the *Causal* condition were more likely to endorse a causal effect than the *Not Causal* condition, $t(41) = 2.52, p = .02$. Taken together, we find that, on average, participants’ *quantitative* judgments varied as a function of data on close alternatives.

However it was also critical to examine participants’ *explanations* to confirm that they were indeed privileging close alternatives (in line with the predictions of Account 2). This was particularly important for understanding the results of the *Causal* condition, as privileging either the data from close alternatives *or* prior causal knowledge would *both* result in stronger causal judgments.

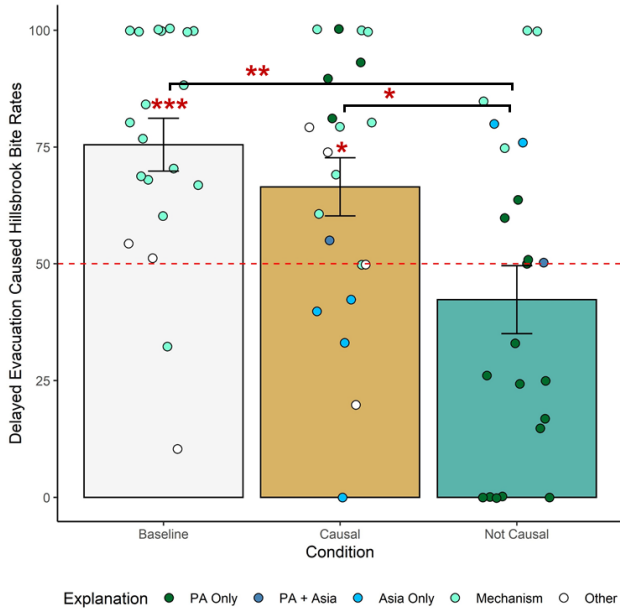


Figure 2: Causal judgments and explanations by condition. Higher scores indicate participants more strongly endorsed that the delayed evacuation *caused* high bite rates. Bars represent standard mean errors. Asterisks above the *brackets* indicate significant condition differences. Asterisks above the *standard error bars* indicate that the mean is significantly different from the midpoint, 50. * $p < .05$. ** $p < .01$. *** $p < .001$.

Explanations

When examining participants' explanations, we considered their relation to condition and causal judgments. All explanations are color coded in Figure 2 and described in greater detail in Tables 1 and 2. As summarized in Table 1, explanations were coded into one of five categories: (1) *PA Only*: participants only mentioned the other Pennsylvania towns' data, and did not discuss the Asian towns ($\kappa = .96$), (2) *PA + Asia*: participants referenced both data from other Pennsylvania towns and Asian towns ($\kappa = 1$), (3) *Asia Only*: participants only mentioned Asian towns' data ($\kappa = .90$), (4) *Mechanism*: participants drew on prior knowledge and discussed the effectiveness of evacuations in general or in this specific case ($\kappa = .94$), or (5) *Other*: all other responses ($\kappa = .86$). Explanations were coded by two independent coders and disagreements were resolved via discussion.

In Table 2, we report the frequency of each explanation type by condition, and the mean causal judgment for each. Like before, higher scores indicate that participants more strongly endorsed that the delayed evacuation was causal.

When participants had *no* additional data beyond the target town (i.e., the *Baseline* condition), they most frequently drew on their prior knowledge and discussed the causal mechanism (86%), and this led to strong causal inferences ($M_{\text{judg}} = 82$).

Table 1: Types of explanations for causal judgments.

Explanation	Definition	Example
PA Only	Only data from other Pennsylvania towns (and not Asian towns) mentioned.	"Looking at the other two towns in Pennsylvania that evacuated had high bite rates even though they evacuated without any delay."
PA + Asia	Data from both Pennsylvania and Asian towns mentioned.	"Because in other cities of Pennsylvania there were less residents bitten. However, in Indonesia and Malaysia there was still a high increase in population being bitten even after being evacuated."
Asia Only	Only data from other Asian towns (and not Pennsylvania towns) mentioned.	"The data for Kapung and Dima suggest that there are external factors that contribute to the population being bitten aside from evacuation time."
Mechanism	Only causal mechanisms about the effectiveness of evacuations in general or in this particular case are mentioned.	"Evacuation prevents people being bitten."
Other	Responses that did not fit any category.	"I think it did."

Table 2: Explanation frequency and mean causal judgment.

Condition	Baseline		Causal		Not Causal	
	%	M_{judg}	%	M_{judg}	%	M_{judg}
PA Only	0%	--	19%	91	68%	24
PA + Asia	0%	--	5%	55	5%	50
Asia Only	0%	--	19%	29	9%	78
Mechanism	86%	82	38%	80	18%	90
Other	14%	38	19%	56	0%	--

Note. M_{judg} = the mean causal judgment for each explanation in the respective condition; higher scores indicate a stronger endorsement of the delayed evacuation as causal.

In the *Not Causal* condition (see the third column of Table 2), participants most frequently referenced the Pennsylvania towns (68%), and this was associated with weak causal judgments ($M_{\text{judg}} = 24$). Notably, however, some participants mentioned Asia (14% in total), and this was associated with *stronger* causal judgments (Asia Only, $M_{\text{judg}} = 78$; PA + Asia, $M_{\text{judg}} = 50$). There were also participants who privileged their mechanism knowledge (18%), which also was associated with stronger causal inferences ($M_{\text{judg}} = 90$).

With respect to the *Causal* condition, recall that participants' quantitative judgments could indicate either that they were privileging the Pennsylvania towns' data *or* their prior causal knowledge. Explanation data revealed that the most frequent explanation referenced prior knowledge (38%; $M_{\text{judg}} = 80$), and a smaller fraction of people discussed the Pennsylvania towns (19%; $M_{\text{judg}} = 91$). Participants who mentioned the Asian towns (24% in total) had, as would be expected, lower causal judgments relative to the rest of the participants (Asia Only, $M_{\text{judg}} = 29$; PA + Asia, $M_{\text{judg}} = 55$).

Taken together, participants' explanations provide support for a more nuanced conclusion, merging aspects of Accounts 1 and 2. That is, people tend to privilege their prior causal mechanistic knowledge when they do not have any additional data *and* when the data on close alternatives support this knowledge (i.e., in the *Causal* condition). However, when the close alternatives *counter* their prior knowledge (i.e., in the *Not Causal* condition), participants tend to override their existing mechanistic knowledge and make causal judgments based on close alternatives. Beyond these trends, it is also important to note that there was substantial variability in which alternatives people found most relevant. For example, some participants focused on the Asian towns (i.e., far alternatives), and this was associated with causal judgments in the opposite direction of the overall means. Thus, the explanation findings also provide at least some evidence in support of Account 3, which proposes that people may vary in the which counterfactuals they find relevant, leading to opposing causal judgments.

Counterfactual judgments

Our study assumes that people's *counterfactual* judgments underscore their causal judgments. To check this assumption, we examined the correlation between participants' causal judgments and counterfactual judgments (i.e., whether an immediate evacuation would have resulted in at least 5% fewer bites in Hillsbrook). We found that there was a strong correlation between these two judgments, $r = .72$, $p < .001$. Importantly, we found that counterfactual judgments also varied systematically across conditions, $F(2) = 8.07$, $p < .001$, $\eta^2 = .21$, following the same pattern as the causal judgments. Participants were more likely to endorse that the delayed evacuation was difference-making in the *Baseline* condition ($M = 77$) and *Causal* condition ($M = 72$) relative to the *Not Causal* condition ($M = 49$), *Baseline* vs. *Not Causal*: $t(41) = 3.43$, $p = .001$, *Causal* vs. *Not Causal*: $t(41) = 2.75$, $p = .009$. Taken together, counterfactual simulation likely undergirds people's causal judgments about unprecedented events.

Discussion

This study examined how a counterfactual theory of causal judgment can be applied to people's reasoning about unprecedented events. Our data suggests that people tend to privilege their prior causal knowledge when they lack additional data, and when evidence from close alternatives supports this knowledge (in line with Account 1). However, when data from close alternatives *counters* their prior causal knowledge, people privilege this counterevidence over their existing mechanistic beliefs (in line with Account 2). Notably, while these were the average trends, we also found individual variability in whether participants privileged close alternatives, far alternatives, or their prior knowledge, and this was coherently linked to differential causal judgments (in line with Account 3). Taken together, a nuanced account is warranted when describing how people make causal judgments about unprecedented societal events.

By focusing on unprecedented events, our study builds on contemporary counterfactual theory that has largely focused on contexts in which people already have rich causal knowledge. Our results suggest that, without additional data beyond the target outcome, people will draw on the knowledge they have rather than opt to withhold causal judgment. Indeed, most people in the *Baseline* condition discussed the effectiveness of evacuations to support their strong causal inferences. This likely indicates that people simulated what happens when an evacuation is or is not delayed, similar to when they imagine what happens when a ball is or is not present in making physical causal judgments. However, as shown in the results from the *Not Causal* condition, people will reliably override these simulations when close alternatives support the opposite inference (e.g., that an evacuation had no effect). This suggests that counterfactual simulation and subsequent causal judgments about unprecedented events may be much more dynamic due to people's uncertainty, such that these judgments are highly dependent on what evidence is available to them.

Recall that there were also several individuals in the *Not Causal* condition who still privileged their beliefs about the effectiveness of evacuations and inferred a causal effect. One open question is: In what contexts are people *more* confident in their existing generative models than available evidence from close alternatives? Does the strength of these prior beliefs lead them to disregard counterevidence? Future experiments could test this by strengthening people's beliefs about a certain policy's effectiveness, for example, saying that a certain policy has always worked for addressing similar problems in the past. In turn, participants may be willing to privilege prior knowledge over data from close alternatives, treating this disconfirming evidence as irrelevant.

Interestingly, there were also individuals who privileged the *far* alternatives (i.e., data from the Asian towns) in their causal judgments. This raises important questions about what alternatives people find relevant. One likely factor that affects relevance is the perceived *stability* of the causal relationship across contexts (Woodward, 2006, 2010). On the one hand, the effectiveness of a public policy intervention may be

perceived as highly context specific, such that people believe that its impact would vary as a function of societal factors like demography and local resources. Indeed, implementing an evacuation could be much more effective in a society in which people have greater resources (e.g., cars) to follow the order and quickly move away. On the other hand, other individuals may assume that the effectiveness of mandatory evacuations should be generally the same across societies, and thus view any data that suggests otherwise to be important counterevidence.

One unexpected finding was that, despite the *quantitative* judgments in the *Causal* condition suggesting that people privileged the Pennsylvania data, the modal explanation for these judgments focused on the effectiveness of evacuations (as opposed to Pennsylvania). This stood in contrast to the finding that people most *commonly* cited the Pennsylvania data in the *Not Causal* condition. This indicates that people override their prior causal knowledge and privilege data on close alternatives only when it is *disconfirming*. One the other hand, when the close alternatives data is *confirming*, people will instead cite their prior mechanistic knowledge to justify their judgments. This may be because people believe they should cite specific data when making a causal judgment that runs counter to their priors, but that they otherwise prefer broad, mechanistic explanations. As noted above, this preference also supports tenets of mechanistic accounts of causal judgment.

While our study paradigm tests which data people privilege when provided two sets of close and far alternatives, we note that the real world is more complicated in several important ways. First, when reasoning about real-world unprecedented events, there are often *many* close alternatives to choose from that may support *different* conclusions. For example, the United States has been compared to a number of other societies with different outcomes (e.g., Australia, Canada, the United Kingdom). Thus, there may be even *greater* individual variability than documented here when people reason about public policies (Mordecai & Connaughton, 2020). Second, future research should test if certain motivations (e.g., political) and sources (e.g., media) can affect *which* alternatives come to mind and seem most relevant (Tversky & Kahneman, 1973). For example, it is possible that people may privilege certain alternatives that support the causal inferences that align with their worldview, and this may be further exacerbated by partisan media outlets emphasizing these alternatives. Future work should explore the extent to which different counterfactual simulations may be an important source of political polarization on societal issues (Markman et al., 2008).

An open question is the extent to which our findings, that focus on unprecedented *societal* events, may apply to *any* event in which people have limited prior knowledge. For example, might this pattern of results emerge when people make causal judgments about a novel physical object? We suspect that participants will still draw on relevant prior knowledge (e.g., about objects in general), and will privilege close alternatives (e.g., another, similar object) if it counters

their priors. However, there may be important differences between societal events and events that are less politically charged. For example, for societal events, people may be more likely to reject counterevidence from close alternatives if it violates deeply held existing beliefs. Testing the extent to which our findings generalize to any unprecedented event will be an important direction for future research.

In sum, we find that people tend to privilege prior causal knowledge when making judgments about unprecedented societal events, but that they will reliably override this knowledge if data from close alternatives offers counterevidence. However, there is also notable individual variation in which alternatives people find most informative. This research lays the foundation for future work on how judgments about unprecedented societal events may be influenced by the strength of people's prior causal knowledge, exposure to various types of evidence, and motivations to make certain causal claims.

Acknowledgments

The writing of this article was supported by the Eunice Kennedy Shriver National Institute of Child Health & Human Development of the National Institutes of Health under Award Number F32HD098777 awarded to Jamie Amemiya. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health. We are grateful to members of the Early Learning and Cognition Lab for their feedback on earlier drafts of this manuscript.

References

- Beauchamp, Z. (2020). *Canada succeeded on coronavirus where America failed. Why?* Vox. <https://www.vox.com/2020/5/4/21242750/coronavirus-covid-19-united-states-canada-trump-trudeau>
- Berlinger, J. (2020). *Australia's coronavirus epicenter records no new cases as the US and Western Europe struggle to contain the pandemic.* CNN. <https://www.cnn.com/2020/10/26/australia/australia-victoria-covid-19-intl-hnk/index.html>
- Byrne, R. M. J. (2016). Counterfactual thought. *Annual Review of Psychology*, 67(1), 135–157. <https://doi.org/10.1146/annurev-psych-122414-033249>
- Connolly, K. (2020). *Coronavirus: How to tell which countries are coping best with Covid.* BBC News. <https://www.bbc.com/news/world-europe-54391482>
- Gerstenberg, T., Goodman, N., Lagnado, D., & Tenenbaum, J. B. (2021). A counterfactual simulation model of causal judgments for physical events. *Psychological Review*. <https://doi.org/10.31234/osf.io/7zj94>
- Gerstenberg, T., Peterson, M. F., Goodman, N. D., Lagnado, D. A., & Tenenbaum, J. B. (2017). Eye-tracking causality. *Psychological Science*, 28(12), 1731–1744. <https://doi.org/10.1177/0956797617713053>

- Kominsky, J. F., & Phillips, J. (2019). Immoral professors and malfunctioning tools: Counterfactual relevance accounts explain the effect of norm violations on causal selection. *Cognitive Science*, 43(11). <https://doi.org/10.1111/cogs.12792>
- Lewis, D. (1973). Causation. *The Journal of Philosophy*, 70(17), 556. <https://doi.org/10.2307/2025310>
- Mackie, J. L. (1974). *The cement of the universe: A study of causation*. Oxford University Press.
- Markman, K. D., Mizoguchi, N., & McMullen, M. N. (2008). "It would have been worse under Saddam:" Implications of counterfactual thinking for beliefs regarding the ethical treatment of prisoners of war. *Journal of Experimental Social Psychology*, 44(3), 650–654. <https://doi.org/10.1016/j.jesp.2007.03.005>
- Mordecai, M., & Connaughton, A. (2020). *Public opinion about coronavirus is more politically divided in U.S. than in other advanced economies*. Pew Research Center. <https://www.pewresearch.org/fact-tank/2020/10/28/public-opinion-about-coronavirus-is-more-politically-divided-in-u-s-than-in-other-advanced-economies/>
- Nolan, D. (2013). Why historians (and everyone else) should care about counterfactuals. *Philosophical Studies*, 163(2), 317–335. <https://doi.org/10.1007/s11098-011-9817-z>
- Tetlock, P. E., & Belkin, A. (1996). *Counterfactual thought experiments in world politics: Logical, methodological, and psychological perspectives*. Princeton University Press.
- Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, 5(2), 207–232. [https://doi.org/10.1016/0010-0285\(73\)90033-9](https://doi.org/10.1016/0010-0285(73)90033-9)
- Wolff, P., & Thorstad, R. (2017). Force dynamics. In *Oxford Handbook of Causal Reasoning* (pp. 147–167). Oxford University Press.
- Woodward, J. (2005). *Making things happen: A theory of causal explanation*. Oxford University Press.
- Woodward, J. (2006). Sensitive and insensitive causation. *The Philosophical Review*, 115(1), 1–50. <https://doi.org/10.1215/00318108-2005-001>
- Woodward, J. (2010). Causation in biology: Stability, specificity, and the choice of levels of explanation. *Biology & Philosophy*, 25(3), 287–318. <https://doi.org/10.1007/s10539-010-9200-z>