The causal role of counterfactuals in responsibility ascriptions to ignorant agents

Katarina Kovacevic[1] (kovacevic_katarina@phd.ceu.edu), Jonathan Kominsky[1] (kominskyj@ceu.edu), Francesca Bonalumi[1,2] (bonalumif@ceu.edu) & Christophe Heintz[1] (heintzc@ceu.edu)
[1]Department of Cognitive Science, Central European University, Vienna, Austria
[2]TUM, Munich, Germany

Abstract

It is now well-established that counterfactual reasoning takes place when people make moral judgments. Less is known about which counterfactuals lead to stronger moral judgment, especially when judging agents who unknowingly produce negative consequences. We explored the relationship between counterfactual salience and responsibility ascription in two experiments. In Experiment 1, we asked people to produce counterfactual alternatives to a vignette they read spontaneously. We manipulated whether agents who produced harm knew the relevant information beforehand and what the reasons for the possible ignorance were. The counterfactual type that people first came up with (e.g., related to external factors or agent's actions) mediated the relationship between the condition and responsibility ratings. Experiment 2 investigated the causal connection between certain counterfactual types and responsibility ascription. We show that guiding people to consider alternative perpetrator's actions leads to a higher tendency to ascribe responsibility than considering victim's actions.

Keywords: counterfactual reasoning; responsibility ascription; ignorance; epistemic actions

Introduction

When something negative happens, we tend to think of alternative events that would lead to a different outcome, or, in other words, we engage in counterfactual simulations (Kahneman & Tversky, 1982; Byrne, 2016). People engage in counterfactual reasoning when they make causality, responsibility, and blame judgments (Kominsky & Phillips, 2019; Kirfel & Lagnado, 2021; Alicke et al., 2008; Gerstenberg, 2024).

Counterfactual reasoning enables the identification of relevant causes (Gerstenberg, 2024; Kirfel & Lagnado, 2021). People are more likely to identify someone's action as a relevant cause, rather than external factors (Hart & Honore, 1985; Hilton et al., 2010). Moreover, people tend to perceive norm violating events as a relevant cause of harm (Icard, Kominsky, & Knobe, 2017). If two agents bring about a negative outcome, people assign more causality (Kominsky & Phillips, 2019) and blame to the agent who violated some norm (Reuter et al., 2015; Samland & Waldman, 2016). However, this is not the case when agents are not aware that they are violating the norm (Kirfel & Phillips, 2023). Hence, when relevant knowledge is lacking, the effect of norm violation on causality and responsibility judgments diminishes.

Kirfel and Lagnado (2021) offer an account in which people do not only imagine alternative actions, but also alternative epistemic states of others when assigning causality for negative outcomes. When the agent knows the relevant fact, counterfactual interventions, or simulations with alternative variable values, focus on changing the agent's instrumental action. On the other hand, if an agent does not know the relevant fact, counterfactual interventions revolve around changing the agent's epistemic state in the form of "had they known." The authors propose that the strength of the counterfactual dependence between the intervention and an alternative outcome could explain why less causality is ascribed to ignorant agents. Alternative instrumental action directly changes the outcome, while having more knowledge does not mean the outcome would certainly have been different. The later depends on other factors, such as the agent's motivation to act differently (Kirfel & Lagnado, 2021).

On the other hand, ignorance does not always exculpate. One of the factors affecting whether the ignorant agent will be judged is the reason behind the agent's ignorance. Intentionally ignorant agents are found more responsible for the negative consequences brought about than unintentionally ignorant ones (Samland & Waldmann, 2016; Kirfel & Lagnado, 2021; Kirfel, Bunk, Zultan, & Gerstenberg, 2023), and in some cases even as responsible as knowledgeable agents (Kovacevic, Bonalumi, & Heintz, 2024). A possible explanation for the absence or weaker exculpatory effect of ignorance in the case of intentionally ignorant agents may lie in what people infer from the decision to remain ignorant. Kirfel and colleagues suggest that when people make judgments, they consider why agents chose not to know the relevant facts. In the case of intentionally ignorant agents, people infer that agents would have acted the same had they known the facts (Kirfel et al., 2023). Adding to this account, we state that counterfactual thinking enables the identification of a lack of prosocial intentions as the cause of the negative event (Kovacevic et al., 2024). If agents cared enough for others, they would have gotten the relevant information in order to prevent harm.

One line of studies explored which counterfactuals people provide when asked to imagine how the negative outcome could have been avoided (Kirfel & Lagnado, 2021). They found that people produce, among others, counterfactuals related to epistemic factors, such as better availability of the information (e.g., if the email did not go to spam) or actions taken to get informed (e.g., if the person read an email). To some extent, this counterfactual depended on the experimental condition – e.g., whether ignorance was self, or externally inflicted. Nonetheless, the counterfactual type –

In D. Barner, N.R. Bramley, A. Ruggeri and C.M. Walker (Eds.), *Proceedings of the 47th Annual Conference of the Cognitive Science Society*

whether the counterfactual related to external factors that would have led to knowledge or the agent's epistemic actions - predicted the level of ascribed causality only in the externally inflicted ignorance condition.

The effect of epistemic counterfactuals on responsibility judgments is hence not completely clear. The question is whether counterfactuals for epistemic actions have the same effect as those for instrumental actions. Extending the ideas from the previous work, we aimed to investigate whether epistemic counterfactuals work the same way in predicting responsibility judgments as counterfactuals related to instrumental actions.

To better understand the role of counterfactual reasoning in responsibility ascription to ignorant agents, we conducted two experiments using a vignette study approach. Following the spontaneous counterfactual production method (Kahneman & Tversky, 1982; Kirfel & Lagnado, 2021), we enriched it with new experimental conditions of ignorance in which we varied the effort needed to get the relevant information, as well as with the knowledge condition. In addition, we provided a consistent categorization of counterfactual types across all experimental conditions to enable direct comparison among different types of ignorance and knowledge. Moreover, instead of including epistemic states as a counterfactual category, we split it into two separate ones of agent's epistemic actions and external epistemic factors.

In Experiment 1, we investigated the type of counterfactuals that most often come to people's minds when they read stories about agents who remained ignorant for different reasons and the relationship between these counterfactuals and their responsibility judgments.

In Experiment 2, we further explored the causal influence of counterfactuals on responsibility ratings. We investigated whether guiding people to consider counterfactuals related to what the perpetrator or the victim could have done differently affects participants' responsibility assessment of both agents.

Experiment 1: Counterfactual production

One way to understand what counterfactuals people find relevant is to pose an open-ended question after letting them read a story (Kahneman & Tversky, 1982; Kirfel & Lagnado, 2021). We take that the counterfactual that comes to mind first is cognitively easier to produce and more salient than the counterfactuals that follow. We expected different types of counterfactuals to be mentioned with different frequencies depending on the reason behind the agent's ignorance, i.e., on our experimental manipulation. We also hypothesized that responsibility ascription is affected by the counterfactual that first comes to mind. We predicted that perpetrator-related-counterfactuals would be mentioned more often in cases of self-inflicted ignorance, where there was an opportunity to learn. At the same time, external-factor-related counterfactuals (including epistemic external factors) would be more common in the case of externally inflicted ignorance. In addition, we predicted that perpetrator-related-counterfactuals would elicit more responsibility to the agent in question than the victim-actions-related counterfactuals or external-factors-related counterfactuals.

Methods

Experiment 1 hypotheses and analyses were fully preregistered https://osf.io/zxqkt. Methods used were in accordance with the international ethical requirements of psychological research and were approved by the Psychological Research Ethics Board (PREBO, Ref: 2023-46) from the Central European University in Vienna for conducting the study.

Participants We recruited 300 participants from Prolific (www.prolific.com) and compensated them at a rate of 9 GBP per hour of participation, evaluated by Prolific as a good payment. All participants were older than 18 and they all gave their informed consent to participate in the study. Selection criteria were that English is a primary, first and fluent language. We excluded participants who did not produce at least three different counterfactuals and did not answer the attention check correctly. The Attention check was a multiple-choice question (e.g. "Which appointment was Rose rushing to?"). We also excluded answers to the counterfactual question that were unrelated to the story, a repetition of the previous answer, and those that were not comprehensible. After applying all the exclusion criteria, we were left with 70 data points in the No knowledge condition, 73 in the Hard effort condition, 72 in the Easy effort, and 69 in the Knowledge condition.

Materials We used 12 vignettes describing real-life social situations. The vignettes for all conditions are taken from Kovacevic and colleagues (2024). The scenarios followed a similar structure: Background information describing the context with main character A doing B; information about epistemic states and actions that we manipulated depending on the experimental condition; main character A doing S; and finally, outcome. We created three different scenarios for every condition. The four experimental between-subjects conditions were 1) No knowledge: The perpetrator does not know the relevant information because that information is not available; 2) Hard effort condition: The cost of acquiring the information is high, and the perpetrator chooses to stay ignorant; 3) Easy effort condition: The cost of acquiring the information is low, and perpetrator chooses to stay ignorant; 4) Knowledge: Perpetrator knows the relevant information. In all conditions perpetrators proceed with the action and produce certain negative consequences. Here is an example of one scenario across the four different conditions below:

Background: Rose was on her way to her dentist appointment at 11 am. She went by car and parked in the only available spot.

No knowledge: The sign saying that spot is reserved was moved by someone. She didn't know about the reservation and she parked there.

Hard effort: She saw there was a sign in front of the parking spot, saying that you should check with the parking attendant whether the spot is reserved. The attendant was on a lunch break, so to check when exactly the spot is reserved she would have needed to wait for him to come back at some point. She didn't do so; she parked there and went to her dentist appointment.

Easy effort: She saw there was a sign in front of the parking spot, saying whether and at what time of the day the spot is reserved. To check when exactly the spot is reserved, she would have needed to read the sign. She didn't do so; she parked there and went to her dentist appointment.

Knowledge: She saw the sign saying the spot was reserved at that time. She parked there.

Outcome: The sign indicates that the spot is reserved between 10am and 2pm for persons with disabilities. Consequently, a man couldn't park anywhere and missed his weekly therapy.

Procedure and design We implemented our stimuli using the Qualtrics software (2022). Participants were instructed to read the stories described and answer a few questions. Each participant was randomly assigned to read one vignette. After reading the story, the participants were prompted to answer the following question: "How do you think this outcome could have been avoided? Please complete the sentence below: "If only....". Participants gave their answers in the essay box underneath the incomplete statement. This question was repeated three more times in a similar fashion: "Now we would like you to do this for the second time and come up with something different," "Please complete this sentence again the third time," and "Please complete this sentence one last time."

Participants' sentences were coded according to the following scheme: a) External (non-epistemic) factors - factors not directly related to the agents' actions or thoughts (e.g "there were more parking spots"; b) External epistemic factors - better availability of the information that is not the consequence of agent's actions, but the environment (e.g. "there was a parking sign"); c) Perpetrator's epistemic actions - actions that lead to learning (e.g. "she read the sign"); d) Perpetrator instrumental actions - actions not related to getting informed (e.g. "she didn't park there"); e) Victim actions - actions that victim could have taken to prevent the negative outcome (e.g. "the man left earlier for his therapy"). In this context, we define a victim as the person in the scenario who faces the negative consequences of other people's actions. Categories were a priori determined and adjusted after a pilot study. For more details on coding scheme and plan, consult the preregistration https://osf.io/zxqkt.

After answering all four open-ended questions, we again showed participants the same vignette before asking them to ascribe responsibility to the main agent. For example, "To what extent do you agree with the statement: 'Rose is responsible for the man missing his therapy'?". Participants were instructed to give answers on a 5-point Likert scale where 1 was "Strongly disagree" and 5 "Strongly agree."

In total, we had four dependent variables: I) Overall counterfactual type frequency: how often is a particular counterfactual category mentioned across all four questions; II) Counterfactual saliency: which counterfactual was the first that participant produced; III) Position of a counterfactual: when the counterfactual type is mentioned, out of four iterations (1, 2, 3, 4); IV) Responsibility rating.

Results and Discussion for Experiment 1

Condition and the salient counterfactual type The relationship between condition and the salient counterfactual type was tested by employing a 4x5 Chi-square test of independence, with the first mentioned counterfactual and condition as variables. Results have shown a significant relationship between condition and counterfactual saliency, $\chi^2(12) = 232$, $p < .001$, Cramer's $V = 0.52$. In the No knowledge condition, people most often considered External epistemic factors (45,7%); in Hard effort Perpetrator's instrumental actions (37%), while in the Easy effort, they most often mentioned Perpetrator's epistemic actions (86%). In the Knowledge condition, Perpetrator's instrumental actions were by far the most common answer (81,2%).

We ran three more Chi-square tests to see whether there was a change in this pattern across counterfactuals on different positions. There was a significant relationship between the condition and the second counterfactual $\chi2(12) = 35.1$, $p < .001$, Cramer's $V = 0.20$, the third counterfactual $\chi2(12) = 28.8$, $p < .01$, Cramer's $V = 0.19$, as well as with the fourth counterfactual $\chi2(12) = 27.2$, $p < .01$, Cramer's $V = 0.18$. However, the effect size was lower than in the case of the first counterfactual.

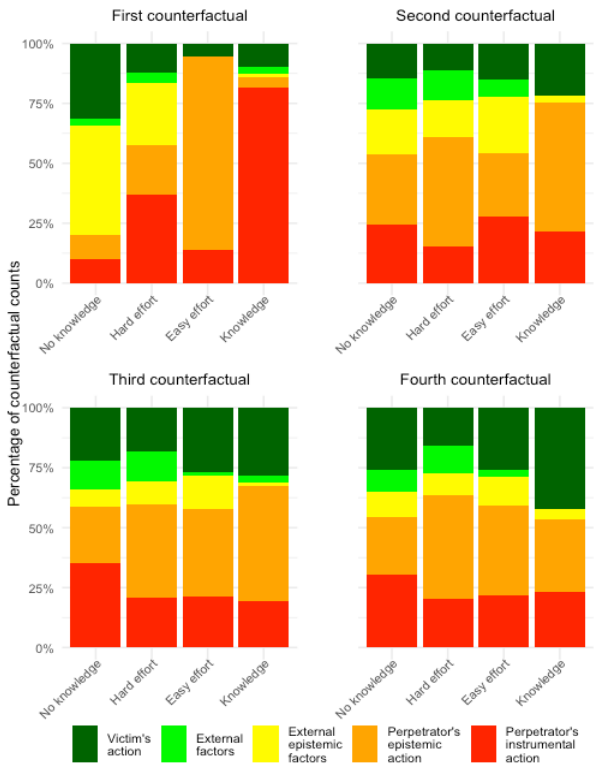

Figure 1: Frequency of certain counterfactual categories by condition and iteration.

As we can see in Figure 1, the mentioned counterfactual type changes with every iteration. The most noticeable differences are the changes in the Knowledge condition – from the perpetrator's instrumental actions to external and victim-related-counterfactuals and the decrease of the perpetrator's epistemic action-related-counterfactuals in the Easy effort condition.

Responsibility and level of knowledge To test the relationship between condition and responsibility ratings, we employed cumulative link mixed models. The tested model included the main effect of the condition and the random effect of the scenario nested on the condition. The model showed a better fit of the data compared to a null model that included only the random effect, $\chi 2(3) = 13.25$, $p < .01$. The results of pairwise comparisons (with p-values adjusted with Bonferroni correction) indicated that the *No knowledge* condition differs significantly from the *Hard effort* condition ($p < .01$) and *Hard effort* condition differs from the *Easy effort* ($p < .001$). In contrast, *Easy effort* and *Knowledge* do not differ ($p = 1$). These results show that ignorance does not excuse the agents if there is an easy opportunity for them to get informed. Furthermore, the effort needed for the epistemic action is considered. Hence, agents who need to put a lot of effort into getting informed are, to some extent, excused for remaining ignorant. These findings fully replicated the results of Kovacevic and colleagues (2024).

Counterfactual saliency and responsibility ascription Figure 2 below shows the relationship between the counterfactual saliency, condition, and responsibility ratings. We see that, independently of the condition, counterfactuals related to the perpetrator's actions (epistemic and instrumental) are most common when participants provide high responsibility ratings. In contrast, victim-related and external (epistemic and instrumental) counterfactuals are most common when participants provide lower responsibility ratings.

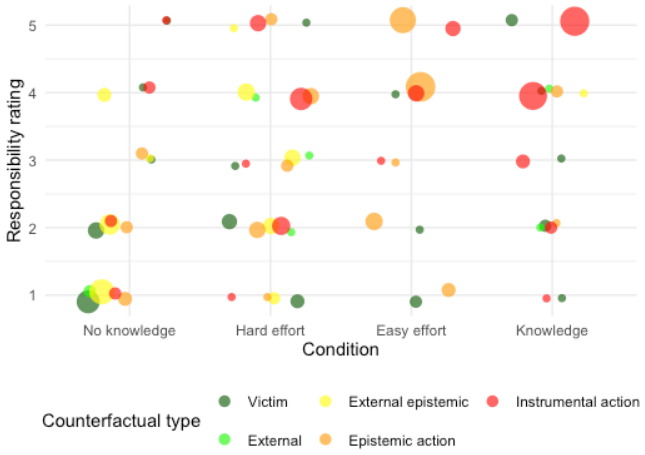

Figure 2: Responsibility ratings by condition and counterfactual type. Each column is a different vignette condition; the responsibility rating is on the Y-axis; the color of each dot indicates the type of counterfactual, and the size indicates frequency of that type of counterfactual.

We also conducted a mediation analysis with a condition as a predictor, counterfactual type as a mediator, and responsibility rating as an outcome. The analysis showed a direct effect of condition on the counterfactual type ($b = 0.570$, $SE = 0.080$, $z = 7.107$, $p < .001$) and a direct effect on responsibility ($b = 0.520$, $SE = 0.067$, $z = 7.728$, $p < .001$). The relationship between counterfactual type and responsibility rating was also significant ($b = 0.281$, $SE = 0.041$, $z = 6.926$, $p < .001$). In addition, an indirect effect of condition on responsibility ratings through counterfactual type was significant, too ($b = 0.160$, $SE = 0.032$, $z = 5.067$, $p < .001$). The total effect of the condition on responsibility ratings was also significant ($b = 0.680$, $SE = 0.061$, $z = 11.091$, $p < .001$). These results suggest that the effect of condition on response is partially mediated by counterfactual type.

Discussion Experiment 1 suggested that the counterfactual type influences responsibility ratings. The pattern indicated that when people first considered victim and externally related counterfactuals, they were less likely to ascribe

responsibility to the main agent than when they first considered the agents' actions. It also showed that people tend to mention different counterfactual categories first than in other iterations, although this may have been influenced to an extent by asking participants to come up with a different answer in every next step.

Experiment 2: Counterfactual manipulation

Since the mediation model in Experiment 1 suggested a causal relationship between counterfactual type and responsibility rating, we aimed to test this causal relationship further. We used the qualitative data from Experiment 1 to create counterfactual prompts that would guide participants to think of a particular counterfactual type.

Prompting people to consider one counterfactual alternative over the other, or in other words, making one alternative more salient, can influence causality judgments. In one study, authors have shown that in the case of two agents equally contributing to an unfavourable outcome, making one's action more salient led to higher causation ascribed to that agent (Phillips et al., 2015). In another study, authors prompted people to consider how the outcome would have been avoided had the agent acted differently or had the machine worked better (Kominsky & Phillips, 2019). Participants ascribed less causality to the agent when they focused on the machine.

Extending the previous studies, we implemented scenarios in which one agent has a role of the perpetrator, and the other of a victim. This experiment hence focused on two counterfactual types – the perpetrator's epistemic actions and the victim's actions. We expect that guiding people to consider counterfactuals related to the victim's actions will lead to a lower level of perpetrator's responsibility than guiding people to focus on what perpetrators could have done better to inform themselves about some relevant aspects of the situation.

Methods

Experiment 2 hypotheses and analyses were fully preregistered https://osf.io/wcavh. Methods used were in accordance with the international ethical requirements of psychological research and were approved by the Psychological Research Ethics Board (PREBO, Ref: 2023-46) from the Central European University in Vienna for conducting the study.

Participants We recruited 400 participants from Prolific (www.prolific.com) and compensated them at a rate of 9 GBP per hour of participation. All participants were older than 18 and they all gave their informed consent to participate in the study. Selection criteria were that English is a primary, first and fluent language. We excluded participants who did not answer the attention check correctly. We also excluded participants who did not provide a comprehensible answer for the "fill in the sentence task". After applying all the exclusion criteria, we were left with 176 data points in the Perpetrator's epistemic actions condition and 170 in the Victim's actions condition.

Materials We used the same scenarios from the previous study, but this time, we kept only one condition – *Hard effort*. This condition was chosen since it showed the most diversity in counterfactual types produced in Experiment 1, suggesting that people's reasoning might be more flexible about this case than others with one highly favored counterfactual type.

Procedure and design We implemented our stimuli using the Qualtrics software. Participants were instructed to read stories described in vignettes and answer a few questions. Each participant was randomly assigned to read one vignette. The two experimental conditions were: 1) Perpetrator's epistemic actions induced by questions such as "If only Ian ran to the bus so he could ask that woman if the money belongs to her, _____, and the woman would have been able to pay her bills that day."; perpetrator's epistemic actions are defined as those actions that agent could have taken to get informed about the relevant aspects of the situation; 2) Victim's actions - induced by questions such as "If only the woman had checked her pockets to see if her money was there, _____, and the woman would have been able to pay her bills that day.".

Participants were instructed to fill in the missing part of the sentence. The idea was that our counterfactual manipulation works stronger if participants actively reason about an alternative causal chain of events.

After reading the story, participants answered two questions that measured the responsibility ascribed to the perpetrator, operationalized with the following question: "To what extent do you agree with the statement: Perpetrator P is responsible for the consequences C"? and the responsibility ascribed to the victim, operationalized with the question: "To what extent do you agree with the statement: Victim V is responsible for the consequences C"? Participants gave answers on a 5-point Likert scale where 1 was "Strongly disagree" and 5 "Strongly agree."

Results and Discussion for Experiment 2

Agent's responsibility and counterfactual type Relationship between the counterfactual condition and perpetrator's responsibility ratings was tested by employing cumulative link mixed models. The tested model included the main effect of the condition, and the random effect of the scenario nested on the condition. The model showed a better fit of the data compared to the null $\chi 2(1) = 6.47$, $p < .05$. Guiding people to consider what actions the victim could have taken led participants to be less likely to ascribe responsibility to the perpetrator (b = -0.81, SE = 0.21, z = -3.80, $p < .001$).

The victim's responsibility and counterfactual type Relationship between the counterfactual condition and the victim's responsibility ratings was tested by employing cumulative link mixed models. The tested model included the main effect of the condition, and the random effect of the scenario nested on the condition. The model did not show a better fit of the data compared to the null, although the p-value was close to the significance level $\chi 2(1) = 3.37$, $p = .067$.

General Discussion

Our first study confirmed that people consider different counterfactuals depending on how much the perpetrator knew before producing negative consequences. When there is no opportunity to get informed, people's counterfactual interventions focus on how the information could have been available and what other agents could have done differently. When the information is available but hard to get, people focus on what the perpetrator could have done differently or how the information could have been more easily available. When the information is easily available, counterfactual interventions focus almost exclusively on what the perpetrators could have done to inform themselves better. Finally, when perpetrators have the relevant knowledge, counterfactual interventions prioritize what they could have done to prevent the harm.

This study also showed that people can come up with multiple counterfactual alternatives to the specific event but that certain counterfactual types are mentioned more often in later instances, some in the first instance. People first consider what the perpetrators could have done differently, including informing themselves, after which their focus goes to external circumstances and reasoning on what the victim could have done to prevent their own harm.

We also found a relationship between the counterfactual type considered and the responsibility ascribed. The reason behind ignorance directly affects how likely responsibility will be ascribed, but the counterfactual type also mediates this relationship. When people first considered counterfactuals related to the perpetrator's epistemic and instrumental actions, they more likely ascribed responsibility to the perpetrator than when considering factors outside the perpetrator's control.

Our second study demonstrated a *causal* relationship between counterfactual salience and responsibility ascription. When people are guided to consider what the perpetrator could have done differently, they tend to ascribe more responsibility than when the victim's alternative preventive actions are made salient. Even though our scenarios included a norm violation, the effect of the counterfactual manipulation is still detectable. Previous studies have shown the effect of counterfactual manipulation in the case of morally neutral scenarios (Phillips et al., 2015; Kominsky & Phillips, 2019). Our results show that people's judgments of responsibility are affected by the information made salient, even in cases of agents violating a social norm, such as taking a parking spot reserved for people with disabilities.

In our second study, we also recorded a tendency of high responsibility ratings ascribed to victims. In one scenario, the victim was even more likely to be perceived as responsible than the perpetrator, with responsibility ratings being the same across both counterfactual conditions. There was a strong normative expectation in this scenario that the victim should be more careful, so the counterfactual manipulation did not affect people's judgments. Interestingly, this scenario also shows us that the ascription of responsibility between the two agents in the story is not a zero-sum game. While the responsibility assigned to the victim did not change in both conditions, the agent in the same scenario received more responsibility in the agent counterfactual condition.

In conclusion, our studies have shown that people consider different counterfactuals when agents produce harm, depending on what agents know before acting. Counterfactuals related to agents' actions lead to more responsibility being ascribed to that agent, while considering external factors or other agents' behavior leads to less responsibility being ascribed. It is also possible to affect people's responsibility assessment to an extent, by making one or the other counterfactual alternative salient.

References

Alicke, M. D., Buckingham, J., Zell, E., & Davis, T. (2008). Culpable control and counterfactual reasoning in the psychology of blame. *Personality and Social Psychology Bulletin, 34*(10), 1371-1381.

Byrne, R. M. (2016). Counterfactual thought. *Annual Review of Psychology*, 67(1), 135-157.

Gerstenberg, T. (2024). Counterfactual simulation in causal cognition. *Trends in Cognitive Sciences.*

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

Kirfel, L., & Lagnado, D. (2021). Changing Minds—Epistemic Interventions in Causal Reasoning.

Kirfel, L., Bunk, X., Zultan, R. I., & Gerstenberg, T. (2023). Father, don't forgive them, for they could have known what they're doing. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 45, No. 45).

Kirfel, L., & Hannikainen, I. R. (2023). Why blame the ostrich? Understanding culpability for willful ignorance. In K., Prochownik, S. Magen,(Eds.), *Advances in experimental philosophy of law,* 75-98.

Kirfel, L., & Phillips, J. (2023). The pervasive impact of ignorance. *Cognition, 231*, 105316.

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), e12792.

Kovacevic, K. M., Bonalumi, F., & Heintz, C. (2024). The importance of epistemic intentions in ascription of responsibility. *Scientific Reports, 14*(1), 1183.

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.

Qualtrics. (2022). *Qualtrics software* [Survey software]. https://www.qualtrics.com

Reuter, K., Kirfel, L., Van Riel, R., & Barlassina, L. (2014). The good, the bad, and the timely: how temporal order and moral judgment influence causal selection. *Frontiers in psychology, 5*, 116135.

Samland, J., & Waldmann, M. R. (2016). How prescriptive norms influence causal inferences. *Cognition, 156*, 164-176.