

# UC Merced

## Proceedings of the Annual Meeting of the Cognitive Science Society

### Title

Revisiting Causal Pluralism: Intention, Process, and Dependency in Cases of Double Prevention

### Permalink

<https://escholarship.org/uc/item/2bn725nx>

### Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 45(45)

### Authors

Thanawala, Huseina  
Erb, Christopher D.

### Publication Date

2023

### Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

# Revisiting Causal Pluralism: Intention, Process, and Dependency in Cases of Double Prevention

Huseina Thanawala (huseina.thanawala24@gmail.com)

Department of Psychology, 23 Symonds Street,  
Auckland CBD, Auckland, 1010, New Zealand

Christopher D. Erb (christopher.erb@auckland.ac.nz)

Department of Psychology, 23 Symonds Street,  
Auckland CBD, Auckland, 1010, New Zealand

## Abstract

Causal pluralism proposes that humans can reason about causes and effects in terms of both dependency and process relations, depending on the scenario. Empirical support for this view is provided by responses to double prevention scenarios in which an actor attempts to bring about an outcome, a preventer attempts to prevent the outcome, and a double preventer intervenes to stop the preventer's prevention attempt. In contrast to the predictions of the causal pluralism account, two pre-registered experiments ( $N_s = 400$  and  $450$ ) indicate (a) that intentional actions are judged to be significantly more causative of an outcome than unintentional actions for both the actor and the double preventer and (b) that reasoners interpret the double preventer's link to the outcome in terms of a process relation. These results underscore the need to revisit fundamental questions regarding how reasoners form and reason over representations of causal scenarios featuring intentional actions.

**Keywords:** causal pluralism; causal reasoning; double prevention; force theory; intentionality

## Introduction

The ability to learn and reason about causal relations is vital to our ability to behave flexibly and adaptively in the world. Given the importance of causal reasoning for such a wide range of human activities, it is unsurprising that the topic has been explored across many fields, including philosophy (Dinh & Danks, 2020), psychology (Alicke et al., 2015; Wolff & Barbey, 2015), economics (Little, 2010), and computer science (Schölkopf, 2022). This widespread interest has resulted in a diverse range of perspectives on how humans reason about cause and effect across different scenarios (Gerstenberg et al., 2021; Lagnado & Channon, 2008; Walsh & Sloman, 2011), age groups (Buchanan & Sobel, 2011; Gopnik et al., 2004), and cultural contexts (Norenzayan & Nisbett, 2000).

Within psychology, two broad theories have been put forward to capture how we represent and reason about causal relations: *dependency theories* and *process theories* (Chang, 2009). Dependency theories link causes and effects by appealing to logical, modal, or statistical relations. A range of dependency theories have been proposed in the literature, including the counterfactual theory of causation (Lewis, 1973), the probabilistic contrast model (Cheng & Novick, 1990), mental model theory (Khemlani, Barbey, & Johnson-

Laird, 2014), and causal model theory (Lagnado et al., 2007; Sloman, Barbey & Hotaling, 2009). In contrast to dependency theories, process theories share the notion that causal reasoning involves the representation of an exchange or transmission of physical quantities between the cause and the effect (Wolff, Barbey & Hausknecht, 2010). Prominent process theories include transference theory (Aronson, 1971; Fair, 1979), conserved quantity theory (Dowe, 1992), and force theory (Wolff, 2014).

Despite the clear differences between the two theoretical approaches, it is plausible that both theories reflect psychologically real and distinct ways of thinking about causal relations. That is, humans may think about causal relations in terms of either dependency or process theories depending on a range of factors, including individual characteristics or aspects of the scenario. Bearing this in mind, Lombrozo (2010) proposed the causal pluralism account which claims that humans can engage in two distinct modes of explanations that are characterized as involving either teleological or mechanistic modes of thinking.

Teleological explanations cite functions or goals, while mechanistic explanations cite causal mechanisms. Goal-directed behavior is often understood in mechanism-independent terms and encourages a dependence perspective of causation, as there is no requirement that a particular method of transmission is involved. By contrast, Lombrozo (2010) proposes that unintentional (accidental) actions are more likely to be understood in terms of mechanistic explanations, which require a transmission or exchange of force to ascribe causality. This implies that unintentional actions support a process perspective on causation.

Lombrozo (2010) suggests that both dependence and process perspectives can influence causal reasoning, but these perspectives can be leveraged differently depending on whether the reasoners adopt a teleological or mechanistic mode of explanation. This general account has been termed causal pluralism, with the central claim being that dependence relations may be weighed more heavily in evaluating causal relations from a teleological mode while process relations may be weighed more heavily in evaluating causal relations from a mechanistic mode.

To test the causal pluralism account, Lombrozo (2010) presented participants with double prevention scenarios featuring three characters: an actor (Alice) who performs an

action that will result in a particular effect, a preventer (Bob) who attempts to prevent the effect from occurring, and a double preventer (Carol) who disrupts the preventer's prevention attempt. Crucially, the intentional status of the actor and double preventer was manipulated across different versions of the scenarios.

Reasoners' judgments of the double prevention scenarios supported the predictions of the causal pluralism account. Consistent with the notion that intentional actions encouraged reasoners to adopt a teleological mode of explanation and, correspondingly, weigh dependence relations more heavily, participants provided similar causal ratings to the actor (A) and the double preventer (C) when they both behaved intentionally (A+C+ condition), despite the fact that one character (the double preventer, Carol) did not share a process relation with the effect (Lombrozo, 2010).

In another condition in which the actions of the actor and double preventer were accidental (A-C- condition), reasoners provided higher causal ratings for the actor than the double preventer. According to causal pluralism, this difference emerged because reasoners were more likely to adopt a mechanistic mode of explanation in the absence of intentional actions by the actor and double preventer, resulting in process relations being weighed more heavily during reasoning. Given (a) that the actor alone shared a process relation with the outcome and (b) that the process relation between the actor and the outcome entails the existence of a dependence relation, the actor was judged as more causal of the outcome than the double preventer by the majority of reasoners.

### Challenges to Causal Pluralism

The current study explores two challenges to the causal pluralism account of reasoning in double prevention scenarios. The first challenge concerns the various ways in which the intentional status of the characters featured in the scenarios could hypothetically impact reasoning. This challenge proposes that intentionality might influence how the causal relations linking the characters to the outcome are formed, rather than how heavily specific relations are weighted during reasoning. The causal pluralism account proposes that intentionality influences causal ascriptions in double prevention scenarios by differentially weighting specific relations during reasoning. This weighting process can be conceptualized as dissociable from the construction of the relations themselves. That is, causal pluralism describes intentionality as influencing the mode of reasoning (teleological vs. mechanistic) that reasoners adopt and, consequently, how heavily specific relations are weighted during reasoning. However, it is possible that the intentional status of the characters also influences how the relations are constructed or represented by reasoners (as opposed to how the relations are weighted during reasoning). For instance, reasoners may represent dependence relations resulting from intentional actions as stronger than those resulting from unintentional actions. Consequently, attempts to test the weighting process proposed by the causal pluralism account

must take into consideration the alternative routes through which intentionality can influence causal ratings.

To address this concern, Lombrozo (2010) conducted an experiment using deviant causal chains in which the actor and double preventer were both described as intending to bring about a particular outcome, but the specific actions taken by the characters that led to the outcome were accidental. Lombrozo (2010) found that the actor received higher causal ratings than the double preventer in scenarios featuring deviant causal links. This finding suggests that describing the double preventer as having the intention to bring about the effect was not sufficient to account for the effects observed in double prevention scenarios featuring non-deviant links.

Although results from the deviant causation conditions lend some support to the notion that intentionality influences the relative weighting of specific causal relations during reasoning, two considerations merit discussion. First, the difference in ratings provided to the actor and double preventer across the deviant and normal double prevention scenarios only reached significance for one of the two conditions tested (Lombrozo, 2010, Experiment 1b), suggesting that intentionality may impact the representation of relations under at least some conditions. Second, results from the normal (non-deviant) condition in which the actor behaved without intention and the double preventer behaved intentionally (the A-C+ condition) revealed descriptively higher causal ratings for the double preventer than the actor (Lombrozo, 2010, Experiment 1a). Given that the actor shared both a process and dependence link with the outcome, it is unclear from the view of causal pluralism why the double preventer would receive higher ratings than the actor unless intentionality influenced the formation or representation of the relation between the double preventer and the outcome. Thus, a primary aim of the current study is to further investigate how reasoning about double prevention scenarios is impacted by manipulating the intentional status of the characters to identify the extent to which intentionality impacts how causal relations are represented (as opposed to how those relations are weighted during reasoning).

The second challenge concerns whether it is indeed the case that reasoners fail to represent a process relation linking the double preventer and the outcome. To motivate this second challenge, we review recent work on force theory, a type of process theory developed by Wolff and Barbey (2015). Force theory particularly emphasizes the role of *allow* and *prevent* relations, though it should be noted that other theories have also developed accounts of these relations (Beller, Bennett & Gerstenberg, 2020; Sloman, Barbey & Hotaling, 2009). Allow and prevent relations are important within force theory for identifying how individual relations can be combined to form causal chains. Such causal chains are then used to obtain new overarching causal relations. Wolff and Barbey (2015) conducted an experiment wherein participants were shown various animations involving car A, car B, and car C. After watching each animation, participants chose the expression that best described the relation between the first and last cars in the chain (e.g., A and C). The results

showed that in a double prevention scenario where A prevents B and B prevents C, 38% of participants answered ‘A caused C’ and 62% of participants answered ‘A allowed C.’

The results from Wolff and Barbey (2015) suggest that participants in Lombrozo’s (2010) study may have represented the double preventer as sharing a process relation with the outcome. Further, it is possible that many of the participants in Lombrozo’s (2010) study believed that the double preventer *allowed* – as opposed to *caused* – the outcome to occur. The current study therefore (a) asks participants to judge how appropriate it was to claim that the outcome occurred *because of* each of the characters’ actions and (b) investigates the frequency with which reasoners judged the actor and double preventer as having *caused*, *allowed*, or *neither caused nor allowed* the outcome to occur.

## The Current Study

The preceding discussion raised three central questions. (1) *To what extent does the intentional status of the actor and double preventer impact how causal relations are represented during reasoning?* (2) *To what extent do reasoners classify the actor and double preventer as having caused, allowed, or neither having caused nor allowed an outcome?* Relatedly, (3) *to what extent are these classifications impacted by the intentional status of the actor and double preventer?* We conducted two pre-registered experiments online to address these questions. The data and analysis files for each experiment are available at [https://osf.io/nreqf/?view\\_only=815a37ead733448d983dd55152032300](https://osf.io/nreqf/?view_only=815a37ead733448d983dd55152032300).

## Experiment 1

Participants in the current experiment were presented with one of the double prevention scenarios developed by Lombrozo (2010) in which three friends (actor = Alice, preventer = Bob, double preventer = Carol) were juggling and listening to music (see Materials). Participants were randomly assigned to one of four conditions: one in which both the actor and double preventer behaved intentionally (A+C+), one in which both behaved unintentionally, (A-C-), one in which only the actor behaved intentionally (A+C-), or one in which only the double preventer behaved intentionally (A-C+). We started the questionnaire by asking participants three forced-choice questions regarding whether the actor, preventer, and double preventer caused, allowed, or did not cause or allow the outcome to occur.

In Lombrozo’s (2010) study, participants were asked to evaluate causal claims (e.g., “Alice caused the music to play.”) on a 6-point scale of appropriateness. We also asked participants to evaluate causal claims on a 6-point scale of appropriateness. However, given the consequential distinction between cause and allow relations highlighted by force theory (Wolff & Barbey, 2015), we avoided the word “cause” in our appropriateness rating statements. Instead, we asked participants to evaluate statements featuring the word “because” (e.g., “The music played because of Alice.”).

We pre-registered four central predictions for this experiment: (1) Participants will most frequently select the cause relation to describe Alice’s contribution to the outcome, the neither cause or allow relation to describe Bob’s contribution to the outcome, and the allow relation to describe Carol’s contribution to the outcome. (2) Participants will be more likely to classify Alice’s contribution to the outcome as causing and less likely to classify the Alice’s contribution to the outcome as not causing or allowing when Alice behaved intentionally. Similarly, participants will be more likely to classify Carol’s action as allowing the outcome and less likely to classify the Carol’s action as not causing or allowing the outcome when Carol behaved intentionally. (3) The appropriateness ratings for Alice will be significantly higher when Alice acts intentionally (in the two A+ conditions) relative to when she acts unintentionally (in the two A- conditions). (4) The appropriateness ratings will be significantly higher for Carol than the Alice in the A-C+ condition.

## Method

**Participants** The pre-registration for this experiment is available at <https://osf.io/87dbh>. We used G\*Power to conduct power analyses for this experiment. We sought to have (a) .80 power to detect an effect size of .25 with a one-tailed independent means t-test at the standard .05 alpha error probability and (b) .80 power to detect an effect size of .25 with a one-tailed paired-samples t-test at the standard .05 alpha error probability. Our power analyses indicated that a sample size of approximately 100 participants in each condition would be needed to reach these goals. We therefore sought to have a final sample size of 400 participants. We only included the first 100 participants from each condition who met the criteria.

We recruited 460 participants (age range 18–65) from the Prolific online participant recruitment platform. Participants received GBP 1.34 in compensation for their time. Participants were excluded from the study if they (a) had not answered one or more of the questions or (b) answered more than one of the comprehension-check questions incorrectly. After participants were excluded according to the aforementioned criteria, we had a final sample with data from 100 participants in each of the four experimental conditions.

**Materials** The experiment consisted of an online survey that presented participants with the double prevention scenario from Lombrozo (2010) provided below involving the characters Alice, Bob, and Carol. An example of one version of the scenario where both the actor (Alice) and double preventer (Carol) acted intentionally is stated below (Lombrozo, 2010):

‘Alice, Bob, and Carol have spent the afternoon juggling and listening to music. At the moment, Alice is juggling, and the music is not playing. Alice wants to listen to music, so she deliberately throws a juggling ball, which heads straight for the stereo’s ‘on’ button. But while Alice’s ball is in the air, Bob starts pulling on the power cord connecting the stereo to

the outlet. If Bob unplugs the cord, it will prevent Alice's ball from turning on the stereo and starting the music. However, Carol wants the music to play, so she deliberately steps on the power cord just before Alice's ball hits the 'on' button, preventing Bob's pull from unplugging the stereo. As a result of these events, the music starts to play.'

Participants in the current study were randomly allocated to one of four versions of the scenario (A+C+, A-C-, A+C- and A-C+). After reading the scenario, participants were asked to respond to questions like the following for each character (example for Alice shown):

Which of the following statements do you agree with the most?

- 1) Alice caused the music to play.
- 2) Alice allowed the music to play.
- 3) Alice did not cause or allow the music to play.

Following these questions, participants had to evaluate the appropriateness of claims regarding Alice, Bob, and Carol on a 6-point scale (from "Completely Inappropriate" to "Completely Appropriate"). Participants judged the appropriateness of a series of sentences stating that the outcome occurred "because of" each character in the scenario (e.g., "The music played because of Alice."). Lastly, we ended the questionnaire by asking participants five true-or-false comprehension questions about the scenario.

**Procedure** The experiment consisted of a questionnaire created on an online platform called Qualtrics. Participants were presented with a description of the study and consented to participate by clicking a link from Prolific that redirected them to the Qualtrics questionnaire. Participants were randomly assigned to one of the four conditions described above. Additionally, the order of the force-choice questions, the order of the statements about each character for the appropriateness ratings, and the order of the comprehension questions were randomized for each participant. The participants could read the scenario at the top of the page when completing the force-choice questions and the questions regarding the appropriateness ratings for each character. For the last section featuring the comprehension questions, participants could not read the scenario on their screen when answering the questions.

## Results

We pre-registered four central predictions for this experiment. Our first hypothesis stated that participants will most frequently select the cause relation to describe Alice's contribution to the outcome, the neither cause or allow relation to describe the Bob character's contribution to the outcome, and the allow relation to describe Carol's contribution to the outcome. Responses to the agreement questions were analyzed with a series of chi-square tests, following Wolff and Barbey (2015). A one-sample chi-square test was performed to examine the causal ratings for each character across all four conditions. The results for Alice yielded a significant effect,  $\chi^2(2, N = 400) = 382.66, p < .001$ , with the cause relation selected significantly more often than

the allow relation or the neither cause or allow relation. The results for Bob yielded a significant effect,  $\chi^2(2, N = 400) = 679.42, p < .001$ , with the neither cause or allow relation selected significantly more often than the cause relation or the allow relation. The results for Carol yielded a significant effect,  $\chi^2(2, N = 400) = 149.43, p < .001$ , with the allow relation selected significantly more often than the cause relation or the neither cause or allow relation.

Our second hypothesis stated that participants would be more likely to classify Alice's action as causing the outcome and less likely to classify Alice's action as not causing or allowing the outcome when Alice behaved intentionally (as opposed to unintentionally). We evaluated how the intentional status of the actor (i.e., Alice) impacted responses to the agreement question for the actor. A chi-square test of independence was performed to examine the relation between the intentionality of Alice and the proportion of each response type. The relation between these variables was not significant,  $\chi^2(2, N = 400) = 2.96, p = .23$ . There was no significant difference in the frequency of cause, allow, or neither responses when Alice behaved intentionally versus unintentionally.

Our second hypothesis also stated that participants would be more likely to classify Carol's contribution to the outcome as "allowing" the outcome and less likely to classify Carol's contribution to the outcome as "not causing or allowing" when Carol behaved intentionally (as opposed to unintentionally). We evaluated how the intentional status of double preventer (i.e., Carol) impacted responses to the agreement question for the double preventer. A chi-square test of independence was performed to examine the relation between the intentionality of Carol and the proportion of each response type. The effect of intentional status had a significant effect on the distribution of responses,  $\chi^2(2, N = 400) = 39.65, p < .001$ . The results showed that the allow relation was selected for Carol significantly more frequently than the cause relation or the neither cause or allow relation when Carol behaved intentionally as compared to unintentionally.

Our next hypothesis stated that the appropriateness ratings for Alice will be significantly higher when Alice acts intentionally (in the two A+ conditions) relative to when Alice acts unintentionally (in the two A- conditions). The average appropriateness judgments given to the actor and double preventer in each condition are presented in Table 1. The appropriateness judgments were analyzed using an independent samples t-test for Alice to see if there is a difference in the mean ratings between the conditions where Alice behaved intentionally (A+C+, A+C-) and unintentionally (A-C+, A-C-). A significant difference was observed between the mean appropriateness ratings for Alice based on intentional status,  $t(398) = -2.27, p = .012$ , with higher appropriateness ratings when Alice behaved intentionally ( $M = 4.84, SD = 1.09$ ) than when Alice behaved unintentionally ( $M = 4.57, SD = 1.23$ ).

Our last hypothesis stated that the appropriateness ratings would be significantly higher for Carol than Alice in the A-

C+ condition. The appropriateness judgments were compared for Alice and Carol using a paired samples t-test in the A-C+ condition only. There was not a significant difference in the appropriateness ratings for the Alice ( $M = 4.56$ ,  $SD = 1.21$ ) and Carol ( $M = 4.74$ ,  $SD = 1.03$ ) conditions;  $t(99) = -1.17$ ,  $p = .12$ .

Table 1: Average Appropriateness Ratings by Condition

Condition	Alice		Carol	
	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>
A+C+	4.85	1.04	4.60	0.98
A-C-	4.58	1.26	3.71	1.31
A+C-	4.82	1.16	4.20	1.05
A-C+	4.56	1.21	4.74	1.03

## Discussion

The results for the first hypothesis showed that participants attributed the cause relation most often to Alice, the allow relation most often to Carol, and the neither cause or allow relation most often to Bob. As predicted, participants were significantly more likely to categorize Carol's behavior as allowing the outcome when she behaved intentionally and significantly more likely to categorize Carol's behavior as neither causing nor allowing the outcome when she behaved unintentionally. This is in line with Wolff's and Barbey's (2015) theory because when participants were given an opportunity to attribute an allow relation to the double preventer, they chose it most frequently. This shows that the double preventer can also be represented as having a process relation to the outcome (in terms of an allow relation instead of a cause relation). Additionally, this shows that although the double preventer is represented as allowing the outcome, these judgements are also influenced by the intentional status of the double preventer.

As predicted, the appropriateness ratings for Alice were significantly higher when Alice acted intentionally (in the two A+ conditions) relative to when Alice acted unintentionally (in the two A- conditions). This indicates that intentionality not only influences how causal relations are weighted during reasoning but also how causal relations are represented. Intentionality per se appears to result in higher appropriateness ratings for the actor, as the weighting process proposed by the causal pluralism account should not result in higher ratings for the actor in the A+ conditions than the A- conditions. However, our results from this experiment did not show a significant difference in the frequency of cause, allow, or neither responses when Alice behaved intentionally versus unintentionally.

Experiment 1 did not reveal a significant difference in the appropriateness ratings for Alice and Carol in the A-C+ condition, though the appropriateness ratings were descriptively higher for Carol than Alice. Given this descriptive difference, along with the descriptive difference reported by Lombrozo (2010, Experiment 1a), we decided to

conduct a second experiment with a larger sample size that focused on the A-C+ condition.

## Experiment 2

In Experiment 1, we hypothesized that the appropriateness ratings would be significantly higher for Carol than Alice in the A-C+ condition. However, the predicted effect did not reach statistical significance. We therefore used G\*Power to determine the sample size necessary to obtain .80 power to detect an effect size of .12 with a one-tailed paired-samples t-test at the standard alpha level of .05. This analysis revealed that a sample of approximately 430 participants was required. We selected a final target sample size of 450 participants.

## Method

**Participants** The pre-registration for this experiment is available at <https://osf.io/78sfj>. Five hundred participants (age range 18–65) were recruited from an online platform (Prolific) and participated in exchange for monetary compensation (GBP 0.75) for a 5-minute study. Participants were excluded from the study if they (a) had not answered one or more of the questions or (b) answered more than one of the comprehension-check questions incorrectly. This resulted in 451 participants who met our inclusion criteria. As stipulated in our pre-registration, our final sample was restricted to the first 450 participants who met the aforementioned inclusion criteria.

**Materials** The experiment consisted of an online survey in which participants were presented with the A-C+ condition used in Experiment 1.

**Procedure** The procedure was identical to Experiment 1 with the exception that participants were only presented with the A-C+ condition of the double prevention scenario.

## Results

Appropriateness judgments for the Alice and Carol characters were analyzed with a one-tailed paired-samples t-test. A significant effect was observed,  $t(449) = -3.16$ ,  $p < .001$ , with higher appropriateness ratings for Carol ( $M = 4.90$ ,  $SD = 0.99$ ) than Alice ( $M = 4.68$ ,  $SD = 1.19$ ).

## Discussion

Experiment 2 revealed significantly lower appropriateness ratings for Alice than Carol in the A-C+ condition, consistent with the descriptive difference reported by Lombrozo (2010, Experiment 1a) and the descriptive difference reported in Experiment 1. This finding presents a challenge for the causal pluralism account as it provides evidence that intentionality per se can influence causal ratings in double prevention scenarios. According to the causal pluralism account, the actor and double preventer should receive equivalent ratings in the A-C+ condition. This is because the actor shares both a dependence and a process relation with the outcome (both of which should be heavily weighted during reasoning),

whereas the double preventer only shares a dependence relation. If intentional and unintentional dependence relations were generated and represented similarly, then the actor and double preventer should receive equivalent ratings in the A-C+ condition. However, this was not the case, which suggests that intentionality impacts how causal relations are represented during reasoning, at least to some degree.

## General Discussion

Understanding how we reason about causal relations is essential to understanding a wide range of human behavior. The current study built on foundational work by Lombrozo (2010) to evaluate two challenges to the causal pluralism account of reasoning in double prevention scenarios.

To evaluate the first challenge, we investigated how the intentional status of the actor and double preventer impacted appropriateness ratings for each character. Results from Experiment 1 showed that the actor received significantly higher ratings in conditions in which the actor behaves intentionally (A+ conditions) relative to conditions in which the actor behaves unintentionally (A- conditions). Heightened appropriateness ratings in the A+ conditions appear to have been driven by information about intentional status altering the nature of the representations linking the actor and the outcome, as the weighting process proposed by causal pluralism should not result in higher ratings for the actor in the A+ conditions than the A- conditions.

To further evaluate the first challenge to causal pluralism, we also examined appropriateness ratings in the A-C+ condition. According to causal pluralism, the actor and double preventer should receive equal ratings in the A-C+ condition because the actor shares both a dependence and a process relation with the outcome (both of which should be heavily weighted during reasoning according to causal pluralism), whereas the double preventer only shares a dependence relation. However, our results from Experiment 2 show that the double preventer receives significantly higher ratings than the actor in this condition which suggests that intentionality does have an impact on how causal relations are represented. Although Experiment 1 revealed descriptively higher ratings for the double preventer than the actor, this difference did not reach statistical significance.

The second challenge raised in the current experiment concerned the causal pluralism account's claim that double preventers are not linked to the outcome via a process link. Lombrozo (2010) presented evidence that reasoners provide relatively low ratings to unintentional double preventers when judging a cause relation. However, research from Wolff and Barbey (2015) indicates that double preventers are often judged to bring about outcomes through allow (as opposed to cause) relations, as predicted by force theory. This finding indicates that a direct transference of force from the double preventer to the outcome is not necessary for reasoners to form a process link between the double preventer and the outcome. This result suggests that the explanatory power of process theories may have been underestimated by the causal pluralism account. Further, this finding raises important

questions regarding the extent to which reasoners maintain multiple, distinct representations when reasoning about the link(s) connecting a character and an outcome. Although it is theoretically possible that reasoners maintain separate dependence and process representations linking characters to outcomes, it is arguably more parsimonious to propose that reasoners instead maintain individual composite representations that reflect a range of factors, including information about dependency, process, and intentionality.

Finally, another important question addressed in this study concerned the way in which the presence or absence of intentional actions affects reasoners' judgments of cause and allow relations for the double preventer. To investigate the extent to which the intentional status of the double preventer impacted these judgments, we directly compared the responses to the agreement questions for Carol when she behaved intentionally versus unintentionally in Experiment 1. Participants were significantly more likely to classify Carol's action as allowing the outcome and significantly less likely to classify Carol's action as not causing or allowing the outcome when Carol behaved intentionally (as opposed to unintentionally). The finding demonstrates that intentional status does have an effect on the allow relations that participants attribute to the double preventer.

## Conclusion

The current study explored two challenges to the causal pluralism account of reasoning in double prevention scenarios. The first challenge concerned the various ways in which the intentional status of the characters featured in the scenarios could hypothetically impact reasoning. Our results showed that intentionality per se does impact the causal ratings attributed to the actor and double preventer, suggesting that intentionality impacts how causal relations are formed as opposed to (or in addition to) how they are weighted during reasoning. The second challenge was inspired by Wolff and Barbey's (2015) research on force theory which indicates that reasoners are quite comfortable describing the link between a double preventer and an outcome in terms of process relations. Our results showed that when reasoners were given the opportunity to attribute an allow relation to the double preventer, they frequently did so, and this attribution was also influenced by the intentional status of the double preventer. These findings indicate that process theories such as force theory have broader explanatory power than proposed by the causal pluralism account. Further, our results highlight fundamental questions concerning whether individuals form, maintain, and reason over (a) separable representations of dependence and process relations or (b) composite representations that reflect information about dependency, process, and intentionality. We believe that revisiting these questions will have important implications not only for our understanding of everyday causal reasoning in humans but also for efforts to design and test artificial systems tasked with learning and interacting with complex causal structures (Glymour et al., 2014; Pearl, 2018; Pearl & Mackenzie, 2018).

## References

- Alicke, M. D., Mandel, D. R., Hilton, D. J., Gerstenberg, T., & Lagnado, D. A. (2015). Causal conceptions in social explanation and moral evaluation: A historical tour. *Perspectives on Psychological Science*, 10(6), 790-812. <https://doi.org/10.1177/1745691615601888>
- Aronson, J. L. (1971). On the grammar of 'cause.' *Synthese*, 22(3/4), 414-430.
- Beller, A., Bennett, E., & Gerstenberg, T. (2020). The language of causation. In S. Denison, M. Mack, Y. Xu, & B. C. Armstrong (Eds.), *Proceedings of the 42nd Annual Conference of the Cognitive Science Society* (pp. 3133-3139). Cognitive Science Society.
- Buchanan, D. W., & Sobel, D. M. (2011). Mechanism-based causal reasoning in young children. *Child Development*, 82(6), 2053-2066. <https://doi.org/10.1111/j.1467-8624.2011.01646.x>
- Chang, W. (2009). Connecting counterfactual and physical causation. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 31, No. 31). <https://doi.org/10.1037/0022-3514.58.4.545>
- Cheng, P. W., & Novick, L. R. (1990). A probabilistic contrast model of causal induction. *Journal of Personality and Social Psychology*, 58(4), 545-567. <https://doi.org/10.1037/0022-3514.58.4.545>
- Dinh, P. P., & Danks, D. (2020). *Causal pluralism in philosophy: Empirical challenges and alternative proposals* [Carnegie Mellon University]. The Philosophy of Science Association. <https://doi.org/10.1086/714878>.
- Dowe, P. (1992). Wesley Salmon's process theory of causality and the conserved quantity theory. *Philosophy of Science*, 59(2), 195-216. <https://doi.org/10.1086/289662>
- Fair, D. (1979). Causation and the flow of energy. *Erkenntnis*, 14(3), 219-250. <https://doi.org/10.1007/BF00174894>
- Gerstenberg, T., Goodman, N. D., Lagnado, D. A., & Tenenbaum, J. B. (2021). A counterfactual simulation model of causal judgments for physical events. *Psychological Review*, 128(5), 936-975.
- Glymour, C., Scheines, R., Spirtes, P., & Kelly, K. (2014). *Discovering causal structure: Artificial intelligence, philosophy of science, and statistical modeling*. Academic Press.
- Gopnik, A., Glymour, C., Sobel, D. M., Schulz, L. E., Kushnir, T., & Danks, D. (2004). A theory of causal learning in children: causal maps and Bayes nets. *Psychological review*, 111(1), 3.
- Khemlani, S. S., Barbey, A. K., & Johnson-Laird, P. N. (2014). Causal reasoning with mental models. *Frontiers in Human Neuroscience*, 8(849), 1-15. <https://doi.org/10.3389/fnhum.2014.00849>
- Lagnado, D. A., & Channon, S. (2008). Judgments of cause and blame: The effects of intentionality and foreseeability. *Cognition*, 108(3), 754-770.
- Lagnado, D. A., Waldmann, M. R., Hagmayer, Y., & Sloman, S. A. (2007). Beyond covariation. In A. Gopnik & L. E. Schulz (Eds.), *Causal learning: Psychology, philosophy, and computation* (154-172). Oxford University Press.
- Lewis, D. (1973). *Counterfactuals*. Blackwell Publishers.
- Little, D. (2010). *Current Issues in the Philosophy of Economics*. Colgate University.
- Lombrozo, T. (2010). Causal-explanatory pluralism: How intentions, functions, and mechanisms influence causal ascriptions. *Cognitive Psychology*, 61(4), 303-332. <https://doi.org/10.1016/j.cogpsych.2010.05.002>
- Norenzayan, A., & Nisbett, R. E. (2000). Culture and causal cognition. *Current Directions in Psychological Science*, 9(4), 132-135. <https://doi.org/10.1111/1467-8721.00077>
- Pearl, J. (2018). Theoretical impediments to machine learning with seven sparks from the causal revolution. *arXiv preprint arXiv:1801.04016*.
- Pearl, J., & Mackenzie, D. (2018). *The book of why: The new science of cause and effect*. Basic Books.
- Schölkopf, B. (2022). Causality for machine learning. In H. Geffner, R. Dechter, & J. Halpern (Eds.), *Probabilistic and Causal Inference: The Works of Judea Pearl* (pp. 765-804). ACM Books.
- Sloman, S., 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. <https://doi.org/10.1111/j.1551-6709.2008.01002.x>
- Walsh, C. R., & Sloman, S. A. (2011). The meaning of cause and prevent: The role of causal mechanism. *Mind & Language*, 26(1), 21-52. <https://doi.org/10.1111/j.1468-0017.2010.01409.x>
- Wolff, P. (2014). Causal pluralism and force dynamics. In B. Copley & F. Martin, *Causation in grammatical structures* (pp. 100-119). <https://doi.org/10.1093/acprof:oso/9780199672073.001.0001>
- Wolff, P., & Barbey, A. K. (2015). Causal reasoning with forces. *Frontiers in Human Neuroscience*, 9, 1. <https://doi.org/10.3389/fnhum.2015.00322>
- 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. <https://doi.org/10.1037/a0018129>