Statistical norm effects in causal cognition

Lara Kirfel (lara-christina.kirfel.15@ucl.ac.uk)
David Lagnado (d.lagnado@ucl.ac.uk)

Department of Experimental Psychology, University College London
26 Bedford Way, London WC1H 0AP England

Abstract

Current causal theories argue that the statistical normality or abnormality of an action makes a difference to people's causal judgements. In this paper, we present two experiments that explore the role of statistical norms in causal cognition. In our first experiment, we provide a preliminary test of two competing theories that aim to explain the effects of normality in causal cognition – the actual causal strength measure (Icard Kominsky & Knobe, 2017) and the correspondence hypothesis about causal judgements (Harinen, 2017). In addition, we control for an often neglected factor, the epistemic states of agents. Our second experiment investigates the effect of statistical normality in the same context, but with a probabilistic rather than deterministic causal structure. Our results favour Icard et al.'s (2017) model of causal strength, but show that the statistical normality of an action loses its influence when the occurrence of the outcome is probabilistic. We discuss the implications of our findings for current causal theories.

Keywords: Causal judgement, statistical norms, normality

Normality in causal cognition

The current literature on causal cognition suggests that the extent to which an action is judged to cause an outcome is influenced by how normal or abnormal this action is regarded. Studies show that actions defined as abnormal, e.g. actions deviating from statistical or prescriptive norms, receive different causal attributions compared to normal actions (Halpern & Hitchcock, 2015; Icard, Kominsky, & Knobe, 2017; Kahneman & Miller, 1986; Knobe & Fraser, 2008). Based on these findings, current theorists have begun to consider the relevance of the normality of actions and events for causal frameworks. (Knobe & Fraser, 2008; Halpern & Hitchcock, 2015; Icard et al., 2017).

A causal structure in which an outcome depends on the occurrence of at least two actions is known as a *conjunctive causal structure*. Studies have shown that when two agents bring about an outcome in a conjunctive structure, higher causal ratings are given to the causal candidate who behaves 'abnormally', e.g. by violating a rule, or performing an infrequent or atypical action. The increased causal attribution to an abnormally acting agent has also been described as the 'abnormal inflation' effect (Icard et al., 2017). Icard, Kominsky and Knobe (2017) tested the influence of prescriptive and descriptive normality on causal judgements in conjunctive as well as disjunctive causal structures, and found a reversed effect of normality in disjunctive cases. When the causal structure is *disjunctive*, i.e. such that the outcome occurs so long as one agent acts, and two agents overdetermine the outcome, the abnormal action is seen as less causal than the normal action. ('abnormal deflation', Icard et al., 2017). These findings show that the direction of the effect of normality on causal attribution depends on the underlying causal structure. While an abnormal causal candidate receives increased causal ratings in a conjunctive causal structure, it receives decreased causal ratings in a disjunctive case.

Modelling normality

Icard et al. (2017) propose a novel *measure of actual causal strength* that aims to predict the patterns of normality in causal judgement as a whole. In their model, the actual causal strength combines two distinct causal measures. The first, *actual necessity* is defined as the counterfactual dependence of the outcome E on the causal factor c (Halpern & Pearl, 2001; Lewis, 2001):

"Actual Necessity":
If C had not occurred, E would not have occurred.

The second, *sufficiency*, defines the robustness with which the causal factor C brings about the outcome E in different contexts:

"Robust Sufficiency":
Given that C occurred, E would have occurred even if the background conditions had been slightly different.

In their model, the strength of a causal factor C is the weighted sum of its necessity and sufficiency, with necessity being weighted by the probability of C being absent, 1 - $P(C{=}1)$, and sufficiency being weighted by the probability of C being present $P(C{=}1)$. Crucially, Icard et al. (2017) suggest that normality directly influences the propensity with which we sample counterfactual situations. The more normal a counterfactual scenario, the more likely it is to be imagined. Given the different weightings of sufficiency and necessity by the probability of a factor being present (respectively absent), the normality of a causal factor will have a different influence on the sampling probability in each case. A sufficient causal factor will be more likely to be imagined when it is normal. In contrast, a necessary factor will be less likely to be imagined when it is normal.

A causal candidate has different degrees of causal necessity and sufficiency for an outcome depending on whether a conjunctive or disjunctive causal structure is in place. By weighting these two causal measures differently by the normality of a cause, the actual causal strength measure allows to predict the overall pattern of normality in conjunctive and disjunctive structures (Icard et al., 2017).

Normality Correspondence

Harinen (2017) has suggested an alternative theory for the different directions of influence of norms on causal judgements - the *correspondence hypothesis about judgement of actual causation* (Harinen, 2017 Gavanski & Wells, 1989; Kahneman & Miller, 1986). He argues that in addition to considering the 'normality' of a causal factor, current accounts also need to take into account the normality of the effect of an action. According to the correspondence theory, abnormal effects, i.e. effects that are statistically or prescriptively abnormal, are also judged to have statistically or prescriptively abnormal causes. Likewise, normal effects, are judged to have normal causes.

Crucially, the correspondence hypothesis provides an alternative explanation for abnormal inflation and deflation with statistical normality. Depending on the kind of causal structure, the frequency or infrequency of an action will also make a difference for the frequency of the outcome. This can be shown if we assume a standard scenario in which outcome E is brought about by agent C_1's action A_1 and C_2's action A_2. C_1 performs A_1 frequently, while C_2 performs A_2 infrequently. Whether E occurs frequently or infrequently will depend on the underlying causal structure. In a conjunctive causal structure, it can be assumed that E occurs infrequently, given that it takes the action of both C_1 and C_2 to bring about E. According to Harinen (2017), in correspondence with the infrequently occurring E, the infrequently acting C_2 will be seen as more causal for E than C_1. In contrast, in a disjunctive structure where E occurs as long as either C_1 or C_2 act, it can be assumed that E occurs frequently, given that C_1 performs A_1 frequently. Corresponding to the frequent occurrence of E, the frequently acting causal agent C_1 will be seen as more causal for E than C_2.

Causal Structure and Knowledge

Both accounts predict why a statistically 'abnormal' or infrequent agent is judged less causal to an outcome in a conjunctive structure, but more causal in a disjunctive structure. While Harinen (2017) refers to the correspondence between the statistical normality of the effect and the statistical normality of the cause, the actual causal strength measure by Icard et al. (2017) predicts the effect based upon the statistical normality of the causal candidate alone. Consequently, one way to distinguish between both theories is to test the effects of abnormality in conjunctive and disjunctive structures while holding the normality of the effect constant. Neither Icard et al.'s (2017) nor Harinen's (2017) account make any reference to the epistemic status of the agents. Whether the agents have any knowledge about the link between their actions and the outcome is not considered in the current literature on norm effects in causal cognition. However, studies have shown that the extent to which agents foresee the consequences of their actions has a crucial impact on causal attributions (Lagnado & Channon, 2008). In particular, higher foreseeability leads to higher causal ratings. Despite these findings, studies investigating the influence of the normality have not systematically controlled for the influence of an agent's epistemic state on causal judgments.

In addition, the influence of prescriptive and statistical norms on causal judgement has mainly investigated actions in deterministic causal structures, i.e. structures in which the effect definitely occurs once the causal conditions are satisfied. However, most of the cause-effect relationships we encounter in our daily lives are probabilistic. We were therefore interested whether the effects predicted by Icard et al. (2017) hold for cases of probabilistic causation.

Consequently, this paper had three aims. First, to provide a preliminary test to distinguish between Icard et al.'s actual causal strength measure (2017) and Harinen's correspondence theory (2017) for statistical norm effects. Second, to control for the influence of an agent's epistemic state on the effects of statistical normality. Third, to test the influence of statistical normality in a probabilistic causal structure.

Experiment 1

The first experiment investigates the influence of statistical norms on causal judgements in two different causal structures. In addition, we varied what agents know about the potential consequence of their actions. In order to test Harinen's (2017) norm correspondence theory of causal judgement, we used a causal structure in which the 'statistical normality', i.e. the frequency of the outcome is fixed.

Participants

600 participants were recruited for this online study via Amazon Mechanical Turk. They were paid £0.50 upon completion of the study. 147 participants were excluded for answering one or more of the three manipulation check questions wrong, leaving a final sample of 453 (215 male, 237 female, 1 N.A.) participants, aged 18-75 years (M = 38.28, SD = 12.62).

Design

We investigated the influence of three factors on causal ratings in a two-agent-scenario: the statistical normality of an action (*frequent* vs. *infrequent action*), the type of causal structure (*conjunctive* vs. *disjunctive*), and the knowledge of the agents about the outcome of their actions (*knowledge* vs. *no knowledge*). Statistical normality, i.e. frequency of actions was manipulated for one agent ('varied agent') while holding the frequency of actions fixed on the second agent ('fixed agent'). The manipulation of statistical normality, causal structure and knowledge was manipulated between-participant.

Material

Participants were presented with a short story in two animated video clips. The videos were programmed using moovly.com. The first video clip introduced the scenario (example clip: https://youtu.be/tS5D-e1uylI). The town "Smallville" has a market place with two central parking spaces for customers. Two agents, called "Ben" ('varied agent') and "Tom" ('fixed agent'), make use of these central parking spaces. However, there is a third agent, a delivery truck driver, who also needs to use these parking spaces in order to deliver products to the adjacent shops. If the delivery truck is prevented from parking there, the delivery cannot be accomplished (Figure 2).

Outcome structure. The outcome in this scenario was defined as the problems that arise when the truck is prevented from parking. In this case, the delivery cannot proceed. We therefore implemented a causal structure in which the outcome is dependent on a causal condition, i.e. the parking spaces being blocked, and an enabling factor, the truck. In our scenario, the delivery truck only comes on Fridays. Thus, the possibility of the outcome to occur is restricted to one day only, given that the causal condition are satisfied, i.e. the relevant parking space is occupied.

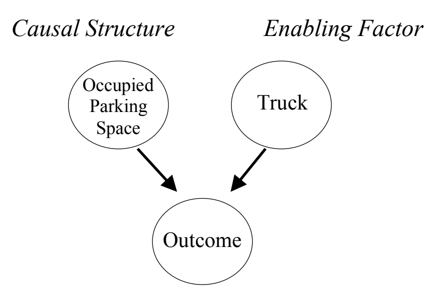

Figure 1. Causal Structure Graph.

Causal Structure The causal structure is manipulated by varying how much space is required for the delivery truck to park (Figure 3). In the '*conjunctive structure*', the truck needs only one of the two central parking spaces, and is blocked when both of them are occupied. In the '*disjunctive structure*', the truck needs both central parking spaces in order to park, and is blocked when at least one parking space is occupied.

Conjunctive Case: $\text{Outcome} = (\text{Tom} \wedge \text{Ben}) \wedge \text{Truck}$
Disjunctive Case: $\text{Outcome} = (\text{Tom} \vee \text{Ben}) \wedge \text{Truck}$

Epistemic Status In the '*no knowledge condition*', there is no indication on the market place that there is a delivery truck using the central parking spaces on Fridays (Figure 4). In the '*knowledge condition*', there is an informatory sign indicating that one (two) parking spaces might be blocked by a delivery truck on Fridays. Importantly, the sign only announces the existence of the delivery truck, but does not state any bans or prohibitions of parking.

Frequency A second video clip shows a sequence of five days, Monday to Friday (https://youtu.be/wIJJ4OJNAaU). The clip varies the frequency with which both agents use the central parking spaces during the week. In the '*Both agents frequent'* condition, both Tom and Ben use the parking spaces from Monday to Thursday, and always arrive at the same time. In the "*One agent frequent, one agent infrequent*' condition, only Tom uses the parking spaces from Monday to Thursday. Finally, in both conditions, Tom and Ben park on the central parking spaces on Friday at the same time. Later, the truck arrives but cannot park. In consequence, the negative outcome was only caused on that last day.

Figure 2. Scenario. General set up of the experimental scenario.

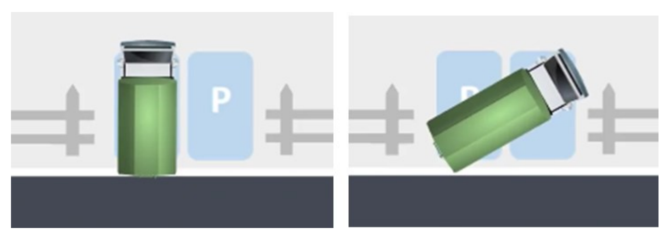

Figure 3. Causal Structure. Left ='conjunctive', right = 'disjunctive'.

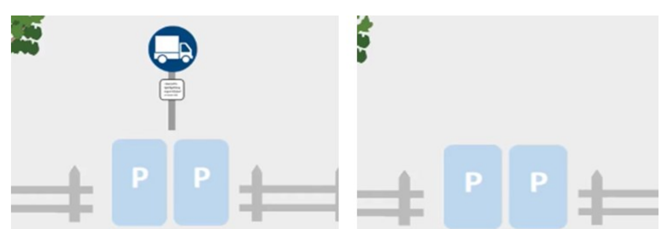

Figure 4. Knowledge. Left ='knowledge', right = 'no knowledge'

Causal Question

After seeing both video clips, the final outcome was again summarised ("On Friday, the delivery truck comes and cannot deliver the new products for the stores"). Participants answered two graded causal judgment questions about the final outcome ["To what extent did Tom (Ben) cause the outcome?"] on an 11 point ["0 – 'Not at all' to 10 - "Fully"] scale.

Results

Causal ratings were analysed with a 2 (causal structure) × 2 (frequency) × 2 (knowledge) × 2 (agent; fixed vs. varied agent) Mixed ANOVA.

The analysis revealed a main effect of causal structure, $F(1, 445) = 23.61$, $p<.001$, $\eta_p^2 =.050$, and knowledge, $F(1, 445) = 45.68$, $p<.001$, $\eta_p^2 = .093$.

Participants generally gave higher causal ratings in the disjunctive structure ($M = 6.52$; $SD = 2.95$) than in the conjunctive structure ($M = 5.27$, $SD = 2.78$).

When the agents knew about the outcome, they were judged more causal ($M = 6.71$; $SD =2.42$) than if they did not know about the outcome ($M = 5.03$; $SD = 3.13$). The analysis also revealed a significant two-way interaction of agent and causal structure $F (1, 445) = 12.09$, $p < .001$, $\eta_p^2 = .026$ and a significant three way interaction between the agent, causal structure and frequency of action $F(1, 445) = 17.11$, $p = .037$, $\eta_p^2 = .037$.

To further explore the three way interaction, we conducted two 2 (frequency) × 2 (agent) ANOVAs for each causal structure condition, i.e. conjunctive and disjunctive.

In the conjunctive causal structure condition (Figure 5), the ANOVA revealed a significant interaction for agent and frequency, $F(1, 234) = 9.15$, $p = .003$, $\eta_p^2 = .038$, and agent, frequency and knowledge, $F(1, 234) = 9.15$, $p = .042$, $\eta_p^2 = .018$. When the agents know about the outcome, the infrequently acting agent ('Ben') is judged to be more causal (M=6.07, SD=2.87) than the frequently acting agent ('Tom'), ($M = 4.86$, $SD = 3.10$), $t(1, 55) = 3.16$, $p = .003$. When the agents do not know about the outcome, there is no difference between an infrequently and frequently acting agent, $t(1, 55) = 0.36$, $p = .155$.

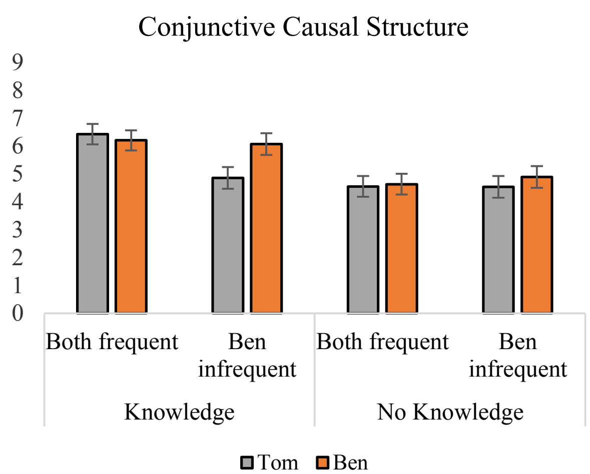

Figure 5. Mean causal ratings (0-10 scale) in 'conjunctive'. Error bars represent ±1 SE mean.

In the disjunctive causal structure condition (Figure 6), the ANOVA revealed a significant interaction for agent and frequency, $F(1, 11) = 10.11$, $p = .002$, $\eta_p^2 = .046$. In the case of Tom acting frequently and Ben acting infrequently, people assign more causal contribution to the frequently acting agent ($M = 6.66$, $SD = 2.88$) than the infrequently acting agent ($M = 6.16$, $SD = 3.10$). A post hoc t-test in the 'knowledge' condition however only revealed a tendency towards significance for the difference between the frequently and infrequently acting agent ($MD = .36$, $SD = 1.41$), $t(1, 55) = 1.90$, $p = .063$.

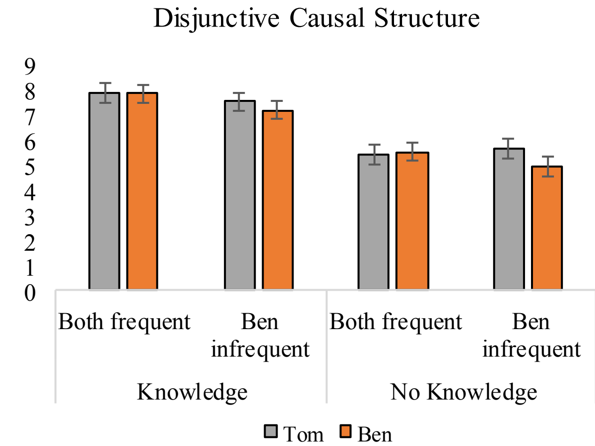

Figure 6. Mean causal ratings (0-10 scale) in 'disjunctive'. Error bars represent ±1 SE mean.

Discussion

Overall, the experiment confirmed the predictions of the actual causal strength model (Icard et al., 2017), showing "abnormal inflation" in conjunctive causal structures and "abnormal deflation" in disjunctive causal structures (Icard et al, 2017). However, by controlling for the epistemic status of the agents, we showed that the abnormal inflation effect only occurs when both agents know about the occurrence of the third agent, the delivery truck driver. When the agents are ignorant of the delivery truck, there is no difference in causal judgements between the frequently and the infrequently acting agent.

Our experiment also showed that causal attributions were generally higher in a disjunctive causal structure than in a conjunctive structure. This finding presents an interesting challenge for current counterfactual theories of causation that predict higher causal attribution for conjunctive than (overdetermined) disjunctive causes (Lagnado, Gerstenberg, & Zultan, 2013; Zultan, Gerstenberg, & Lagnado, 2012).

In line with Lagnado and Channon (2008), we found that causal attributions are higher when the agents know that the delivery truck comes on Friday, and thus the potential outcome is more foreseeable.

Crucially, our experiment confirms the predictions of Icard et al. (2017) while holding the frequency of the effect constant by fixing it to one particular day ('Friday'). Our findings therefore distinguish between the actual causation theories of Harinen (2017) and Icard et al. (2017). Harinen's correspondence theory (2017) holds that the underlying psychological mechanism of the effects described by Icard et al. (2017) is linked to the correspondence between the normality of cause and effect. The results showed that "abnormal inflation" of causal attribution in conjunctive structures and "abnormal deflation" in disjunctive structures occur without co-variation of the statistical frequency of cause and effect.

Experiment 2

The aim of the second experiment was to investigate the influence of frequent vs. infrequent actions on causal judgements in a probabilistic causal structure. Applied to our experimental scenario, we explored whether statistical normality still influences causal judgements when the occurrence of the third agent (the truck) on a particular day is not certain.

Participants

950 participants were recruited for this online study via Amazon Mechanical Turk. They were paid £0.50 upon completion of the study. 215 participants were excluded for not answering one or more of the four manipulation check questions wrong, leaving a final sample of 735 (309 male, 422 female) participants, aged 18 -79 years (*M* = 39.43; *SD* = 13.14).

Design

The experiment investigated the influence of four factors on causal judgements in a two-agent-scenario: frequency (*frequent* vs. *infrequent action*), causal structure (*conjunctive* vs. *disjunctive*), knowledge (*knowledge* vs. no *knowledge*), and probability of the outcome (*high* vs. *low*).

Probability In the second experiment, the occurrence of the truck on Friday was probabilistic. There was no particular day on which the truck would carry out the delivery, instead the truck could come on any day with a particular frequency. The information about the probability with which the delivery could occur on a certain day was varied in the introduction clip. In the '*low probability*' condition, the delivery truck comes once every 4 weeks, in the '*high probability condition*', the truck comes once every 9 days (example clip: https://youtu.be/fA7wwWXH17M). The sequence of days presented was as in Experiment 1. Again, both agents park on Friday and the truck arrives and cannot park. However, the occurrence of the truck on that day was three times more likely in the high probability condition compared to the low probability condition.

Results

The analysis of causal judgements revealed significant main effects of causal structure $F(1, 719) = 6.22$, $p = .013$, $\eta_p^2 = .009$, and knowledge $F(1, 719) = 12.01$, $p = .002$, $\eta_p^2 = .016$. There was a significant interaction between causal structure and knowledge $F(1, 719) = 4.58$, $p = .030$, $\eta_p^2 = .007$. The results did not yield a significant interaction between 'agent' and 'frequency', $F(1, 719) = 0.04$, $p = .85$. nor 'agent', 'frequency' and 'knowledge' $F(1, 719) = 0.13$, $p = .72$, indicating that there were no differences in causal judgements between the agent whose statistical normality was varied, and the one whose normality was fixed.

We explored the interaction between causal structure and knowledge in two separate 2 (causal structure) × 2 (knowledge) × 2 (Agent) ANOVAS for each knowledge condition. When both agents have knowledge about the third agent, there is a significant difference between the conjunctive ($M = 4.85$, $SD = 2.66$) and disjunctive structure ($M = 5.92$, $SD = 3.27$), $F(1, 344) = 10.99$, $p = .001$, $\eta_p^2 = .031$. However, when both agents do not know about the possible outcome of their actions, there is no significant difference between causal attributions in the conjunctive ($M = 4.56$, $SD = 2.94$) and the disjunctive structure ($M = 4.65$, $SD = 3.23$), $F(1, 375) = .05$, $p = .821$ (Figure 7).

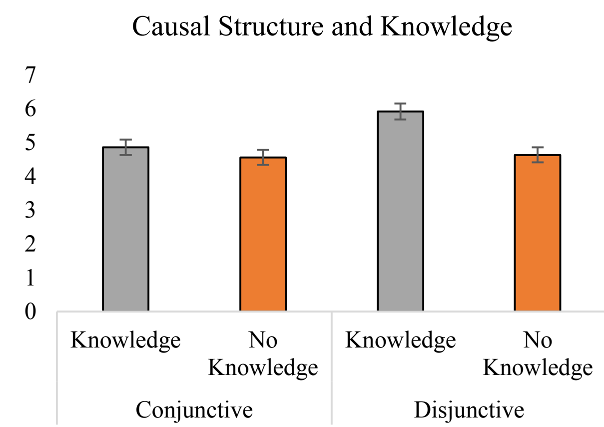

Figure 7. Mean causal ratings (0-10 scale) by causal structure and epistemic status condition. Error bars represent ±1 SE mean.

Discussion

Our experiment showed that the causal structure and the knowledge of the agents have an impact on causal ratings. In line with our findings from Experiment 1, participants attributed higher causality to agents when the agents had knowledge about the probabilistic occurrence of the truck. However, an interaction revealed that this difference is only prevalent in the disjunctive structure, while knowledge does not make a difference when the underlying causal structure is conjunctive. Crucially, we did not find evidence for the influence of the statistical normality of an agent's action in a probabilistic causal structure.

General Discussion

In our two experiments we have explored the effect of statistical normality on causal judgements. In Experiment 1, we found evidence for the two effects predicted by the actual causal strength measure proposed by Icard et al. (2017). However, we showed that abnormal inflation in a conjunctive causal structure disappears when the agents have no knowledge about the outcome. In contrast, abnormal deflation is weakened when the agents know about the outcome. These findings show that the epistemic status of agents can be a crucial factor for the influence of statistical normality. While in a disjunctive structure, the outcome of an action is foreseeable by the agent as soon as they know their action and the causal structure, the foreseeability in a conjunctive structure depends both on the knowledge of their action as well as knowledge about the frequency of the other agent's behaviour. Crucially, in the '*one agent frequent, one infrequent*' condition, our experiment did not specify whether both agents knew about each other's behavioural pattern. Lagnado and Channon (2008) showed that an agent receives higher causal and blame ratings when the outcome of the action is foreseeable.

As a result, our findings raise the question whether the effect of statistical normality on causal judgements is partly based on foreseeability advantages of agents that arise from the statistical normality of their actions. In line with abnormal inflation in conjunctive cases (Icard et al, 2017), a sudden one-off behaviour might lead to higher foreseeability of the outcome in conjunctive causal structures if the one-time acting agent knows that the other agent acts frequently. In a disjunctive structure, the outcome occurs as soon as one agent acts. As a result, in a disjunctive structure, a frequently acting agent will have encountered the consequences of their action multiple times. In line with abnormal deflation in disjunctive structures, this might lead to a knowledge advantage and hence more causal attribution to the frequently acting agent, compared to a first time actor.

In addition, the first experiment attempts to provide a first test between the actual causal strength model by Icard et al (2017) and Harinen's theory (2017). Despite the fixed frequency of occurrence in both causal structures, we demonstrated 'abnormal inflation' and 'abnormal deflation' in causal judgements. As a result, we take this as supportive evidence that the effects of statistical normality predicted by Icard et al. (2017) occur without an additional corresponding normality between cause and effect. By fixing the possibility of the occurrence of the outcome to one particular day, we tested a scenario in which the effect is rare, or statistically abnormal. However, as a consequence our study did not present a situation in which the effect occurs frequently. Future studies will need to systematically vary the statistical normality of effects and actions as well as the underlying causal structure in order to provide a comprehensive test for Harinen's correspondence theory (2017).

Crucially, our experiments have shown that the implementation of a deterministic vs. probabilistic causal structure influences whether the statistical normality of an agent's action will actually exert influence on causal attributions. Our second experiment showed that the statistical normality of an agent's action does not influence causal attributions when the link between actions and outcome is probabilistic. While knowledge and the type of causal structure continued to influence causal attributions, the statistical normality of an action does not have an impact, neither in combination with certain epistemic states or a certain probability of the outcome.

As a result, our study raises the question of why the switch from a deterministic to a probabilistic link between actions and outcome reduces the effect of statistical normality on causal judgements. Again, one possible reason could lie in the difference of epistemic uncertainty in a deterministic vs probabilistic structure. Lagnado and Channon (2008) found causal judgements to be influenced by *subjective foreseeability*, i.e. how likely an event is from the agent's point of view, but also by the *objective foreseeability* of an outcome, i.e. how likely the event is from an objective point of view. Naturally, both objective and subjective foreseeability of an outcome are lower in a probabilistic than in a deterministic causal structure. More research is needed to test the influence of statistical normality in deterministic vs non-deterministic causal structures. Future work might address the question whether a certain degree of foreseeability of the outcome, based upon the frequency of the agents' action, is needed for an influence of statistical normality on causal judgements.

References

Driver, J. (2008). Attributions of causation and moral responsibility. In W. Sinnott-Armstrong (Ed.), *Moral psychology, Vol. 2. The cognitive science of morality: Intuition and diversity* (pp. 423- 439). Cambridge, MA: MIT Press.

Gavanski, I., & Wells, G. L. (1989). Counterfactual processing of normal and exceptional Events. *Journal of Experimental Psychology*, *25*, 314–325.

Halpern, J. Y., & Hitchcock, C. (2015). Graded causation and defaults. *The British Journal for the Philosophy of Science*, *66*(2), 413–457.

Halpern, J. Y., & Pearl, J. (2001). Causes and explanations: A structural-model approach - Part II: Explanations. *IJCAI International Joint Conference on Artificial Intelligence*, *56*, 27–34.

Harinen, T. (2017). Normal causes for normal effects: Reinvigorating the correspondence hypothesis about judgments of actual causation. *Erkenntnis*, 1–22.

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

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

Knobe, J., & Fraser, B. (2008). Causal judgment and moral judgment: Two experiments 1. *Moral Psychology*, *2*(1992), 441–8.

Hitchcock, C., & Knobe, J. (2009). Cause and norm. *Journal of Philosophy*, 106(11), 587–612.

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

Lewis, D. K. (2001). *Counterfactuals*. Oxford: Blackwell Publishers.

Samland, J., & Waldmann, M. R. (2015). Highlighting the causal meaning of causal test questions in contexts of norm violations. *Proceedings of the 37th Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.

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

Sytsma, J., Livengood, J., & Rose, D. (2012). Two types of typicality: Rethinking the role of statistical typicality in ordinary causal attributions. *Studies in History and Philosophy of Science Part C :Studies in History and Philosophy of Biological and Biomedical Sciences*, *43*(4), 814–820.

Zultan, R., Gerstenberg, T., & Lagnado, D. A. (2012). Finding fault: Causality and counterfactuals in group attributions. *Cognition*, *125*(3), 754–770.