I know what you did last summer (and how often). Epistemic states and statistical normality in causal judgements

Lara Kirfel (ucjulki@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

When several causes contributed to an outcome, we often single out one causal factor as being "more of a cause" than others. What explains this selection? Existing research suggests that people's judgements of actual causation can be influenced by the degree to which they regard certain events as norm-deviant, or "abnormal" (Hart & Honoré, 1963; Kahneman & Miller, 1986; Hitchcock & Knobe, 2009; Halpern & Hitchcock 2015). In this paper, we argue that statistical abnormality influences causal judgements about human agents by changing the agents' epistemic states (*Epistemic Hypothesis*). In Experiment 1, we replicate previous findings that people assign more causal strength to a statistically abnormally acting agent, but show that they also assign them more knowledge about the behaviour of their peers. In Experiment 2, we show that in case of equal epistemic uncertainty, people do not differentiate between statistically abnormal and normal causal agents. In Experiment 3, we explore the difference between type and token abnormality, and find that a token abnormal, but type normal behaviour still influences causal judgments, with people's epistemic judgments mirroring these causal judgments. We discuss the implications of this research for current norm-frameworks in causal cognition.

Keywords: statistical norms, normality, causal judgment, counterfactual reasoning, epistemic states

Our ability to form causal judgements plays a fundamental role in human cognition. In everyday life, we encounter situations that demand an explanation of why something happened, how it happened, or how it could have been prevented. Fortunately, our environment is rich in statistical information. Statistical patterns have been shown to be a reliable cue in guiding people's causal inferences and judgements (Cheng, 1997). The co-variation of cue and outcome, their proximity in space and time or the temporal order in which events occur have been shown to inform assumptions about *causal structure*, i.e. the existence of a causal relation between cue and outcome, as well as *causal strength*, i.e. the degree of a causal relation between cue and outcome (Lagnado, Waldmann, Hagmayer & Sloman, 2007).

Recent research suggests that the influence of statistical information on causal cognition goes even further. The *statistical normality* of a causal factor, i.e. how likely, typical or frequent it is perceived, can make a difference to people's causal judgement about this factor over and beyond its actual causal contribution (Cheng & Novick, 1991; Hitchcock & Knobe, 2009; Samland & Waldman, 2016; Kominsky, Phillips, Gerstenberg, Lagnado & Knobe, 2015; Icard, Kominsky, Knobe, 2017). In a range of empirical studies, people have been shown to differentiate between causal factors according to their statistical features, even when both factors are necessary for the outcome to occur (Hitchcock & Knobe, 2009; Icard et al., 2017; Gerstenberg & Icard, n.d.).

Most prominently, this research suggests that deviations from statistical normality increases the causal strength assigned to a cause. Specifically, people are more inclined to judge that *C* causes *E* when *C* is perceived to be statistically "abnormal", i.e. unlikely, infrequent or atypical manner, rather than when C is perceived to be statistically normal. This holds even when in both cases, *C* is known to have the same actual causal contribution to the effect. These findings raise the question of why people take statistical features into account even when these features do not function as supplementary cues to causal structure or strength. What makes people prefer abnormal causal candidates?

Normality matters – but why?

A prominent line of research argues that norms or normality influence causal judgments by changing the relevance or propensity to consider counterfactual possibilities (Kahneman & Miller, 1986; Hitchcock & Knobe, 2009, Icard et al., 2017). A statistical norm violation increases the likelihood of thinking about an alternative scenario in which the norm-violation is replaced by norm-conforming behaviour. A typical test case in this research is causation in a conjunctive causal structure, where two causes are each necessary to produce an outcome. When both C_{normal} and $C_{abnormal}$ together bring about outcome E, people will be more likely to envisage a counterfactual scenario in which $C_{abnormal}$ is absent, rather than a counterfactual in which C_{normal} is absent. According to the counterfactual account, imagining a counterfactual alternative in which normality, or norm-conformity, is restored highlights the causal role of the abnormal causal factor for the outcome, compared to that of the normal causal factor (Kahneman & Miller, 1986; Hitchcock & Knobe, 2009, Icard et al., 2017).

Counterfactual accounts of norm effects in causal cognition have gained increasing popularity. On the one hand, they have integrated norms into formal causal frameworks that can explain a variety of norm effects on causal judgments, such as "causal superseding" (Kominsky

et al., 2017) or "abnormal deflation" (Icard et al., 2017; 2018). On the other hand, they not only predict the influence of statistical norms on people's causal judgements, but also the impact of other kind of norms, such as prescriptive norms (Hitchock & Knobe, 2008) or norms of proper functioning (Phillips & Kominsky, 2018). Recently, it has been suggested that the influence of both prescriptive and statistical norms on causal judgements can be explained by a single normality concept (Bear & Knobe, 2017).

Knowing me, knowing you

The majority of studies supporting the counterfactual account has been conducted using vignette stories in which participants rate the causal impact of human agents who differ in certain aspects of normality. This has led some to argue that the influence of moral abnormality on causal judgements in the context of human agents reveals something about people's blame responses, rather than a difference in counterfactual and causal reasoning (Samland & Waldman, 2016, Alicke, Rose & Bloom, 2012). However, most research argues that statistical norms influence the underlying process of causal judgement. When it comes to statistical norms, it is the abnormality itself that leads people to judge a causal difference between an abnormally and a normally acting causal agent.

In this paper, we propose an alternative hypothesis. While we agree that statistical likelihoods can have an impact on people's causal judgements about events or objects, we think that in the context of human agents, there is another important factor to consider. Epistemic states, i.e. the knowledge an agent has about their environment, have been shown to influence how we evaluate the causality of their actions (Lagnado & Channon, 2008). Whether an agent engages in a frequent of typical action, or an infrequent or atypical action, will likely change their epistemic states about the consequences of this action. In particular, in the case of conjunctive causal structures, an abnormally acting agent seems to have an epistemic advantage over the normally acting agent in knowing or expecting the outcome to happen. We believe it is the epistemic advantage that arises from a statistically abnormal action, rather than the abnormality per se, that drives the main difference in people's judgements about causal agents. We call this the *Epistemic Hypothesis* (EP). We conducted three experiments to investigate this hypothesis. In Experiment 1, we replicate previous literature by showing that people assign more causal strength to a statistically abnormally acting agent. In Experiment 2, we show that in case of equal epistemic uncertainty, people do not make a causal difference between abnormal and normal causal agents. In Experiment 3, we find that a token abnormal, but type normal behaviour still influences causal judgments, with people's epistemic judgments mirroring these causal judgments. We discuss the implications of this research for current norm-frameworks in causal cognition.

Experiment 1

The term "statistical abnormality" has been used broadly in the causal cognition literature, referring to actions or events that are unlikely, rare or atypical. In our experiments we have concentrated on statistical normality in the sense of the frequency of an action. We follow the current paradigm of assessing causal ratings of two causal agents in a conjunctive causal structure, while varying the statistical normality of their actions. In order to focus our investigation, we deviate from the current experimental paradigms in two aspects. Instead of descriptive vignettes ("Agent X frequently does action Y"), we use sequential animated video scenes in order to represent action frequencies more naturalistically. Furthermore, previous literature has suggested that the co-variation between cause and effect influences causal considerations (Harinen, 2017, Cheng 1997, Kirfel & Lagnado, 2018). Current experimental studies are ambiguous about the statistical normality of the effect, which is why we decided to employ a causal structure which allows us to control the frequency of the outcome.

Participants[1]

176 participants were recruited for this online study via Amazon Mechanical Turk. 10 participants were excluded for answering more than one check question wrong, leaving a final sample of 166 (M_{age} = 37.19, SD_{age}=11.24, *age range*= [20-77]; 101 male, 64 female, 1 N.A.) They were paid £0.70 upon completion of the study (Ø 8.06min).

Design

We manipulated two factors in a two-agent-scenario: the statistical normality of an action (*frequent* vs. *infrequent action*) and the type of scenario (*microwave* vs. *coffee machine*). Statistical normality, i.e. frequency of actions was manipulated for one agent (Agent 2: *varied agent*) while holding the frequency of actions fixed on the second agent (Agent 1: *fixed agent*). The scenario type was manipulated between-participant, while the statistical normality was manipulated within-participant. Participant saw two video clips ("*frequent", "infrequent"*) from one of the two scenario types, presented in randomized order. Names of the agents were varied across all conditions.

Material

The frame story consists of two co-workers in a shared office. Depending on the scenario type, the office has either two coffee machines or two microwaves that the employees can use. For energy saving purposes, the company introduces the "Green Friday" on which the building is switched into a power-saving mode. As a result, the use of more than one coffee machine (microwave) on Fridays will lead to a power failure in the building. All workers are aware of the Green Friday.

[1] The material and data for all experiments are available under: https://osf.io/zhvsb/

Figure 1. Scenario "Coffee Machine" with Fixed Agent ("Henry") and varied Agent ("James").

Response Measures

Causal Rating. After each video clip participants were asked to express their agreement with statements about the causal contribution of each agent to the outcome ["Agent 1 (2) has caused the power failure."] on a 7-point Likert scale [1 – 'Strongly disagree' to 7 – 'strongly agree']. Questions were presented in randomized order.

Manipulation Checks. In two subsequent manipulation check questions, participants were asked about their understanding of the action frequency in the scenario ["Who used a coffee machine frequently (rather than infrequently) this week?" – 'Agent 1', 'Agent 2'; multiple answers possible] and the causal structure ["The use of how many coffee machines does it take to produce a power failure on Friday? – 'One coffee machine', 'Two coffee machines']. At the end of the survey, i.e. after watching both videos and answering the causal rating questions, participants were asked to express their opinion about the epistemic states of the agents in both videos ["Agent 1 (2) knew that Agent 2 (1) would use a microwave on Friday."] on a 7-point Likert scale [1 – 'strongly disagree' to 7 – 'strongly agree']. By this, we wanted to check for people's assumptions of the agent's epistemic states.

Results

A Mixed ANOVA for participant's agreement ratings about the causal statements revealed a significant interaction for Frequency × Agent, $F_{(1,164)} = 29.05$, $p < .001$, $\eta_p^2 = .15$. While people judge no difference between the causal contribution of the agents when both of them have frequently performed the action, an agent whose action is rare is seen as more causal ($M = 5.52$, $SD = 1.63$, 95% CI [5.27, 5.78]) then a frequently acting agent ($M = 4.54$, $SD = 1.97$, 95% CI [4.24, 4.84]).

There was no effect for scenario type ($p = .653$). A Mixed ANOVA for agreement ratings about the agent's epistemic states revealed a significant interaction for Frequency × Agent $F_{(1,164)} = 291.60$, $p < .001$, $\eta_p^2 = .64$. When the two agents differ in the frequency of their actions, people express more agreement with the proposition that the agent acting for the first time on Friday knows that their (frequently acting) coworker would act ($M = 5.83$, $SD = 1.74$, 95% CI [5.54, 6.07]), than vice versa ($M = 2.37$, $SD = 1.54$, 95% CI [2.11, 2.64]).

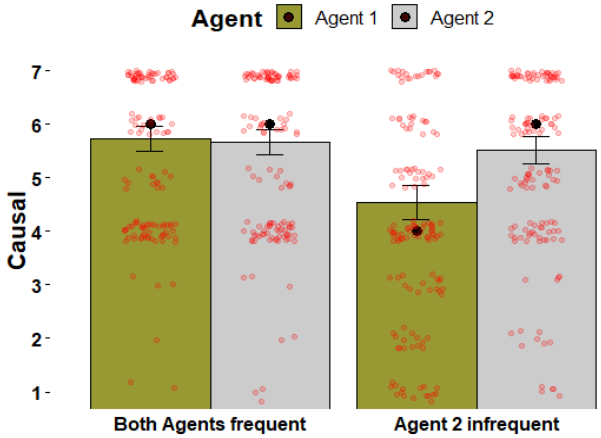

Figure 2. Mean agreement ratings (scale 1-7) for causal statement. Error bars represent ±1 SE mean, black points represent the median.

Discussion

In this experiment we found that when a frequently and infrequently acting agent together cause an outcome, people judge the agent who has acted infrequently to be of greater causal strength than the frequently acting agent. Our findings are in line with the literature in causal cognition showing that people tend to assign more causal strength to abnormal causes (Hitchcock & Knobe, 2009, Icard et al., 2017). In our study, we manipulated the statistical normality among agents' actions. However, in a two-agent conjunctive structure, acting abnormally gives the agent a better chance of foreseeing the consequences of their action. This is because the infrequent worker has witnessed the frequent worker acting on multiple occasions, whereas the frequent worker has never seen the infrequent worker act. In accordance with this prediction, we found that people assigned more knowledge about the co-worker's behaviour to the abnormally acting agent. This leaves open the question whether it was the epistemic advantage of the abnormally acting agent, or the abnormality of their action, that led people to make a causal difference. For our second experiment, we therefore examined whether abnormality still influences causal judgements when there is no such epistemic advantage.

Experiment 2

In the second experiment, we aimed to investigate the effect of statistical normality on causal judgments when neither agent knows about the frequency of the other's actions.

Participants

171 participants were recruited for this online study via MTurk; 19 were excluded for answering more than one check question wrong (N=152, M_{age} = 38.22, SD_{age} =11.25, *age*

range= [19-71]; 81 male, 79 female,). They were paid £0.70 upon completion of the study (Ø 8.61min).

Design & Material

The experiment was designed as Experiment 1, with the difference that the two agents are shown as working in separate offices on different floors. The agents are introduced as co-workers who "[despite] working for the same company, do not know each other and have never met or seen each other." (https://youtu.be/dYaXueuGOoA).

Response Measures

We used the same Causal Rating Measures and Manipulation Checks as in in Experiment 1.

Figure 3. Scenario "Coffee Machine" with fixed Agent 1 ("Henry") and varied Agent 2 ("James").

Results

A Mixed ANOVA for participant's agreement ratings about the causal statements revealed a main effect for Frequency $F(1,150) = 9.96$, $p = .002$, $\eta_p^2 = .06$. Higher causal ratings are given when both agents act frequently ($M = 5.07$, $SD = 1.87$, 95% *CI* [4.87, 5.28]), compared to the case in which only one has acted frequently ($M = 4.72$, $SD = 2.06$, 95% *CI* [4.60, 4.93]).

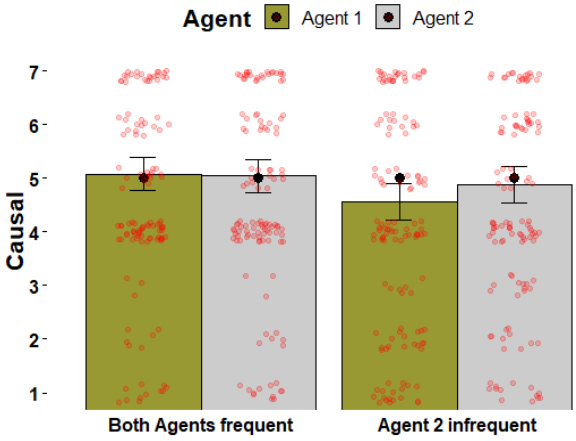

Figure 4. Mean agreement ratings (scale 1-7) for causal statement. Error bars represent ±1 SE mean, black points represent the median.

There was no interaction effect of Frequency × Agent ($p = .118$), and no effect of scenario type ($p = .441$).

A Mixed ANOVA for agreement ratings about the agent's epistemic states revealed a significant interaction for Frequency × Agent $F_{(1,150)} = 4.83$, $p = .029$, $\eta_p^2 = .03$.

Discussion

In our second experiment, we investigated whether statistical normality influences causal judgments when neither agent knows about the other's behaviour. We found that people do not differentiate between a frequently and rarely agent when neither agent knows or observes the other's behaviour. When both agents operate out of sight from each other, people do not judge the abnormally acting agent as contributing more to the joint outcome of their actions. However, the epistemic manipulation check questions revealed that our manipulation of epistemic uncertainty was only partly successful. Although both agents were introduced as working from different offices and not knowing each other, participants still assumed a very small epistemic difference when they differ in their action frequency. Compared to Experiment 1, however, the epistemic difference is negligible ($MD_{EXP1} = 3.36$, $MD_{EXP2} = 0.16$) and rated at the bottom of the 7 point Likert scale [1 – 'strongly disagree' to 7 – 'strongly agree'] (*Frequent Agent*: $M = 1.24$, *infrequent Agent* $M = 1.40$).

As a result, people overall disagreed with the statement that the agents had knowledge of each other. Our experiment shows that the general reduction in the agents' knowledge about each other led to an absence of influence of statistical normality. If an agent has not secured knowledge about the behaviour of their peers, people do not take into account the statistical normality of the agent's behaviour when making causal judgements. Our second experiment therefore shows that in case of epistemic uncertainty, i.e. when acting abnormally does not generate an epistemic advantage, statistical normality does not affect causal judgement.

Type and Token Normality

Our two experiments so far confirm the hypothesis that statistical normality influences causal judgments by giving an epistemic advantage. However, there is another interesting case to consider. Statistical abnormality does not necessarily need to lead to an epistemic advantage when agents, despite differing in their action frequency, can still predict the general outcome-causing behaviour. This case might be hard to experience naturally, because it is exactly the unpredictability of abnormal behaviour that makes it difficult for other agents to foresee it, leading to an epistemic asymmetry. However, when the agent acts for the first time, but their specific action has been performed frequently before by someone else, the agent's behaviour is still abnormal, but others might have been able to foresee the occurrence of this type of action. Strictly speaking, in such a case the abnormality of the behaviour is abnormal only in a limited sense. The agent is abnormal on an "agent-token" level, i.e. *this* particular agent performing action φ, but normal on an "agent-type" level, i.e. *an agent* performing action φ. In their

paper "Two types of typicality", Sytsma, Livengood and Rose (2015) reassess the role of statistical normality by distinguishing between agent-level and population-level statistical norms. They find that agent-level statistical normality has an influence on causal attributions, while deviating from a populational-level norm does not affect people's causal judgements.

For our third experiment, we adopted a similar paradigm as used by Sytsma et al. (2015). We introduced a third 'auxiliary' agent who uses one of the outcome triggering devices regularly during the week before the abnormally acting agent uses it on Friday. By this, we were interested whether an action that is token abnormal, but type normal, still influences causal judgment. Crucially, we assumed that introducing type normality might also make a difference to the agents' epistemic states. That is, in contrast to Experiment 1, here we would expect the token normally acting agent to have certain foreseeability that *someone* performs the causally relevant action on Friday (even though on that day, this happens to be a different agent than expected). The manipulation of epistemic states in Experiment 3 however is much noisy and occurs indirectly through the manipulation of type normality. In line with EP, we predict that if people continue to judge the token abnormal agent to be more causal for the outcome, this would again be tracked by a perceived epistemic asymmetry between these agents.

Experiment 3

In the third experiment, we aimed to investigate the effect of statistical normality on causal judgments when an agent acts statistically abnormal, but their action has been performed before by others.

Participants

180 participants were recruited for this online study via Amazon Mechanical Turk; 26 were excluded for answering more than manipulation wrong (N=154, M_{age} = 38.47, SD_{age} =12.16, *age range* = [19-72]; 90 male, 62 female, 1 2.A). They were paid £0.70 upon completion of the study (Ø 8.64min).

Design & Material

We used the same scenarios as in Experiment 1, but added a third causally irrelevant agent, Agent 3. The statistical normality of the agents who are causing the final outcome was manipulated as before, i.e. varied for one agent and held fixed for the other (*Agent 1*: fixed agent; *Agent 2*: varied agent). In the condition in which both Agent 1 and Agent 2 behave statistically normal, both of them use a coffee machine (microwave) from Monday to Friday, with Agent 3 simply being present and not acting (https://youtu.be/Tsxt1peUA74). In the condition in which Agent 2 acts abnormally, Agent 2 uses the coffee machine (microwave) on Friday, but Agent 3 uses that exact same coffee machine (microwave) the days before, i.e. from Monday to Thursday (https://youtu.be/k2wE52iZPKY).

Response Measures

We used the same Causal Rating Measures as in Experiment 1, but for the sake of completeness, added a Causal Rating for Agent 3 which we did not include in our analysis. We added a Manipulation Check Question to test whether people correctly perceived who had acted on the final day of the outcome ["Who used a microwave on Friday?" 'Agent 1', 'Agent 2', 'Agent 3', multiple answers possible]. We changed our Epistemic Question into a question about i) the type of behaviour "Agent 1 (2) knew that the other coffee machine (microwave) would be used by someone on Friday", and ii) the behaviour of the specific agent "Agent 1 (2) knew that Agent 2 (1) would use the other coffee machine (microwave) on Friday" [1 – 'strongly disagree' to 7 – 'strongly agree'].

Figure 5. Scenario "Coffee Machine" with fixed Agent 1 ("Dan"), varied Agent 2 ("Eddie") and 'auxiliary' Agent 3 ("Sam").

Results

A Mixed ANOVA for participant's agreement ratings about the causal statements about Agent 1 (fixed) and Agent 2 (varied) revealed an interaction effect for Frequency × Agent $F_{(1,152)} = 9.89$, $p = .002$, $\eta_p^2 = .06$.

When Agent 1 and 2 differ in the frequency of the actions that they perform on Friday, people agree more with the statement that the infrequently acting Agent 2 caused the outcome (M = 5.05, SD = 1.99, 95% CI [4.73, 5.36]), than that the frequently acting Agent 1 (M = 4.48, SD = 2.11, 95% CI [4.15, 4.82]).

A Mixed ANOVA for agreement ratings about agent's epistemic states for the type of behaviour revealed a significant interaction for Frequency × Agent $F_{(1,152)} = 10.82$, $p = .001$, $\eta_p^2 = .07$. The infrequently acting agent is judged to have more certainty that *someone* would use the other device on Friday (M = 4.73, SD = 2.01, 95% CI [4.41, 5.05]), than the frequently acting agent (M = 4.23, SD = 2.12, 95% CI [3.89, 4.56]). A Mixed ANOVA for ratings on the agent's assumptions about the *specific agent* using the other device also revealed a significant interaction for Frequency × Agent $F_{(1,152)} = 110.01$, $p < .001$, $\eta_p^2 = .42$. Participants agreed substantially more with the statement that the infrequently acting agent knows that the frequently acting agent would be

using the other relevant device on Friday (M = 4.61, SD = 2.01, 95% CI [4.29, 4.93]), than vice versa (M = 2.53, SD = 1.83, 95% CI [2.24, 2.82]).

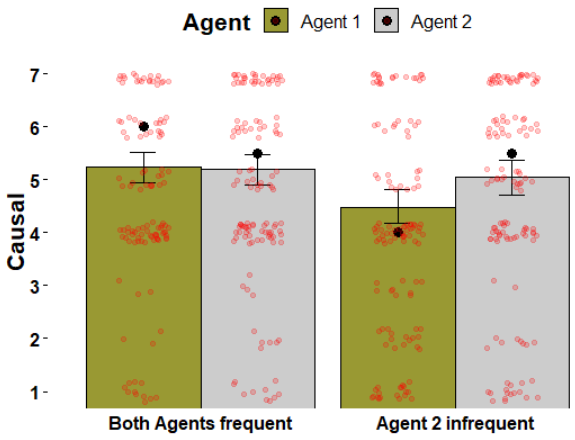

Figure 4. Mean agreement ratings (scale 1-7) for causal statements. Error bars represent ±1 SE mean, black points represent the median.

Subgroup Analysis. We conducted an additional analysis for the causal agreement ratings of the subgroup of people who rated the type behaviour expectations of normal and abnormal agent as equal (n=98). Here, we found no significant interaction for Frequency × Agent ($F_{(1,96)}$ = 1.6, p = .147) ($MD_{\text{Abnormal-Normal}}$=0.29, SD_{MD}=2.0).

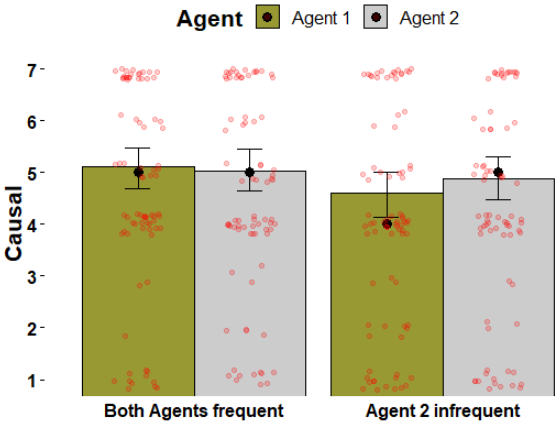

Figure 5. Mean agreement ratings (scale 1-7) for causal statements. Error bars represent ±1 SE mean, black points represent the median.

Discussion

In our third experiment, we found that an action that is token abnormal, but type normal, still influences causal judgments,. However, the judged difference between token normal and abnormal agent is significantly smaller than in Experiment 1. In addition, we found that people thought that the normally acting is less certain that the abnormally acting agent would act, but also less certain that *someone else* would act. This result comes as a surprise, given that both focal agents should have been able to expect *an agent* to act in the final scenario. While we assessed the focal agents' expectations towards the general type and each other's token behaviour, we did not assess their predictions about the behaviour of the third 'auxiliary' agent. It is therefore likely that some people might have assumed Agent 1 (and/or Agent 2) to have expected Agent's 3 omission. In consequence, it might be that the difference in action type expectations comes about as a difference in expectations about who *in fact* acted on Friday. This, again, leaves the normal agent with an epistemic disadvantage. However, a subgroup analysis showed that participants who assumed that both agents had equal behaviour type knowledge, i.e. that both agents were equally expecting that someone would act on Friday, did not judge a significant causal difference between abnormal and normal agent.

General Discussion

In three experiments, we investigated what we call the *Epistemic Hypothesis* (EP), the hypothesis that statistical abnormality will influence causal judgments via generating an epistemic asymmetry. In our first experiment, we showed that an abnormally acting agent is seen as more causally effective for an outcome, but also as more knowing about the behaviour of their normal counterpart. In accordance with EP, we found that in the case of mutual ignorance about each other, statistical abnormality does not influence causal judgements. Finally, we found that token abnormal, but type normal behaviour still influences causal judgments. At the same time, people's epistemic judgments about type and token behaviour mirror these causal judgments.

What role do epistemic states play in the influence of normality on causal judgements? Samland and Waldmann (2016) have shown that the mental states of agents can affect whether people's judgements about their causal contribution are influenced by prescriptive abnormality. They found that people do not take prescriptive norms into account for their causal judgments when the norm-violating agent is unaware of their norm transgression. Counterfactual accounts leave open under which circumstances people start to perceive a behaviour as "abnormal" (Phillips & Kominsky, 2018). Therefore, an agent's lack of knowledge or awareness of existing norms might determine whether the behaviour is perceived as norm-violating or abnormal in the first place. However, we think that in case of statistical normality, an agent can assess the normality status of their behaviour relative to their own action history, their agent-level normality. In consequence, the assessment of statistical normality is not necessarily conditional on the knowledge about external factors, such as rules or laws, or the behaviour of other people. In this paper, we aim to make different claim. We argue that it is the epistemic state that occurs qua the normality or abnormality of an action that drives the difference in people's causal judgements (Kirfel & Lagnado, 2017; Kirfel & Lagnado, 2018). Our experiments support this hypothesis. Hence, we argue that current norm incorporating causal frameworks are in need of a firm theory of epistemic states in order to explain their influence on norm-based causal cognition.

References

Alicke, M. D., Rose, D., & Bloom, D. (2011). Causation, Norm Violation, and Culpable Control. *Journal of Philosophy, 108*(12), 670–696.

Bear, A., & Knobe, J. (2017). Normality: Part descriptive, part prescriptive. Cognition, 167, 25–37.

Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review, 104*(2), 367–405.

Cheng, P.W., & Novick, L.R. (1991). Causes versus enabling conditions. *Cognition*, 40, 83-120

Gerstenberg, T. & Icard, T.F. (submitted). Expectations affect physical causation judgments. Pre-print.

Hagmayer, Y., Sloman, S.A., Lagnado, D.A. and Waldmann, M.R. (2007): Causal Reasoning through Intervention. In: Gopnik, A. and Schulz, L., Eds., *Causal Learning: Psychology, Philosophy, and Computation*, Oxford University Press, Oxford, 86-100.

Halpern, J. & Hitchcock, C. (2015). Graded Causation and Default. *British Journal for the Philosophy of Science*, 66 (2):413-457.

Harinen, T. (2017). Normal Causes for Normal Effects: Reinvigorating the Correspondence Hypothesis About Judgments of Actual Causation. *Erkenntnis 82*(6) 1–22.

Hart, H. L. A., & Honoré, T. (1963). Causation in the Law, 2nd ed., Oxford: The Clarendon Press.

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

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.

Kirfel, L. & Lagnado, D. (2017). Oops, I did it again. The impact of frequency on causal judgements. In G. Gunzelmann, A. Howes, T. Tenbrink, & E. Davelaar (Eds.), *Proceedings of the 39th Annual Conference of the Cognitive Science Society*. (pp. 2420-2425). Austin, TX: Cognitive Science Society

Kirfel, L. & Lagnado, D. (2018). Statistical Norm Effects in Causal Cognition. In T.T. Rogers, M. Rau, X. Zhu, & C. W. Kalish (Eds.), *Proceedings of the 40th Annual Conference of the Cognitive Science Society*. (pp. 615-620). Austin, TX: Cognitive Science Society.

Lagnado, D. A., Waldmann, M. R., Hagmayer, Y., & Sloman, S. A. (2007). Beyond covariation: Cues to causal structure. In A. Gopnik & L. Schulz (Eds.), *Causal learning: Psychology, philosophy, and computation* (pp. 154-172). New York, NY, US: Oxford University Press.

Lagnado, D. A., & Channon, S. (2008). Judgments of cause and blame: The effects of intentionality and foreseeability. *Cognition, 108*, 754–770.

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

Phillips, J. S., & Kominsky, J. F. (2018). Immoral professors and malfunctioning tools: Counterfactual relevance accounts explain the effect of norm violations on causal selection. *Pre-print*.

Kominsky, J. F., Phillips, J., Gerstenberg, T., Lagnado, D., & Knobe, J. (2015). Causal superseding. *Cognition , 37*, 196–209.

Samland, J., & Waldmann, M. R. (2016). How prescriptive norms influence causal inferences. *Cognition, 156*. 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 Biological and Biomedical Sciences*, 43(4), 814-820.