

“Oops, I did it again.”

The impact of frequent behaviour on causal judgement

Lara Kirfel (lara-christina.kirfel.15@ucl.ac.uk)

David Lagnado (d.lagnado@ucl.ac.uk)

Department of Experimental Psychology, UCL, 26 Bedford Way, London WC1E 6BT

Abstract

Current causal theories aim to incorporate the effect of statistical and prescriptive norms on causal judgements, stating that norm-violating actions are judged as more causal than norm-conforming ones. In this paper, we present two experiments that undermine this claim, showing that people attribute increased causality to agents who conform to the norm of frequent behaviour. Furthermore, we find that the time point at which a moral norm is introduced does not make a difference to causal attributions, but that the frequency of a norm violation further accentuates its causal rating. Because these findings present a challenge to current norm theories of causation, we argue for an extended counterfactual model of causal attribution.

Keywords: causal judgement; counterfactual reasoning, frequency; norms; moral judgement

Introduction

Judgements about causal processes are a crucial skill to make sense of the world. While our ordinary concept of causation was long thought to be entirely objective in character, we know that factors that go beyond the causal facts of a situation, e.g. the moral status of an action, are hugely influential on how we think about the causality of an agent (Halpern & Hitchcock, 2015; Hitchcock & Knobe, 2009; Alicke, Rose & Blome, 2011). For example, if two agents cause a negative outcome by jointly performing the same action, but only one of them is actually allowed to perform this action, people judge the agent who violated the rule as more of a cause (Knobe & Fraser, 2008). A similar picture arises for the statistical properties of an action. Exceptional or atypical actions (Kahneman & Miller, 1986), as well as actions with statistically unlikely outcomes (Kominsky et al., 2015) are judged more causal than typical actions or likely outcomes.

All these cases involve circumstances in which the factor that is viewed as a cause is ‘abnormal’: a person does something immoral, a person does something she usually does not do, or an action has a very unlikely outcome. It has therefore been suggested that our causal judgements are sensitive to ‘normality’, or ‘norms’, with the term norms encompassing a variety of norms like moral norms, social rules or statistical norms such frequent behaviour or the likelihood of an outcome. This view holds that norm-violating actions, i.e. actions that deviate from norms or normal circumstances, are judged as more causal than norm-conforming ones (Hart & Honoré, 1963; Hitchcock & Knobe, 2009). The underlying assumption is that we assess a causal candidate in terms of its counterfactual relation for the

outcome (Gerstenberg & Lagnado, 2010; Halpern & Pearl, 2010) and that norms come into play concerning which counterfactuals we consider. Hence, when reasoning about the cause of an event, norms are thought to make norm-conforming counterfactuals more relevant (Hitchcock & Knobe, 2009), more available (Kahneman & Miller, 1986), more probable to be sampled (Icard & Knobe, 2016) or to be ranked highest in the order of possible counterfactuals (Halpern & Hitchcock, 2015).

Common to all these accounts is the idea that in a causal scenario with multiple causal candidates, we will single out the action or event that is norm-deviant, e.g. immoral or unlikely, because we have the tendency to counterfactually simulate an alternative scenario in which normality is in place and no norm is violated (Hitchcock & Knobe, 2009; Kahneman & Miller, 1986). In addition, it has been shown that the norm violation of one agent not only increases causal attribution, it can also reduce causal attribution to other, norm conforming agents, known as “causal superseding” (Kominsky et al., 2015). Although this account applies to normality and norms in a very broad sense, studies in causal attribution have mainly investigated the influence of violations of moral norms (Alicke, Rose & Blome, 2011; Hitchcock & Knobe, 2009; Knobe & Fraser, 2008), statistical norms (Kahneman & Miller, 1986; Kominsky et al., 2015) and norms of proper functioning (Hitchcock & Knobe, 2009). We will subsume these theories under a general norm-framework of causal judgement.

In the causal cognition literature, frequent, repeated actions have often been considered as belonging to a special kind of statistical norms, so called ‘agent-level’ statistical norms or norms of ‘typical behaviour’. According to the norm framework of causal judgement, actions that deviate from this norm, i.e. atypical actions that are rarely or almost never performed, are seen as more causal than actions that are performed frequently (Knobe & Fraser, 2008). However, Sytsma and colleagues showed that these ‘atypical’, i.e. infrequent actions receive reduced instead of increased causal attribution (Roxborough & Cumby 2014; Sytsma et al., 2012). The authors systematically varied the frequency and morality of actions in the ‘pen case’ scenario (Knobe & Fraser, 2008). In this scenario, both a professor and an administrative assistant each take one of two available pens in the department office, with negative consequences. Sytsma et al. varied whether each agent frequently took pens from the office, and whether they were officially allowed to do so. While, as expected, they found that an agent was judged as more causal when violating the department policy, they also

found that an agent was attributed more causality when he frequently took pens rather than only once. This was even the case when there was no department policy prohibiting the use of pens (Sytsma et al., 2012).

By showing that an action that has been repeatedly performed previously receives higher rather than lower causal ratings, the authors provide evidence that agent-level frequent behaviour does not always follow the general schema attributed to the role of norms in causal reasoning. However, the influence of frequent behaviour on causal judgements in this particular case might not come as a surprise. The causal structure of the “pen case” scenario is sensitive towards the frequency of the actions – trivially, the more often a person takes pens, the more they contribute to the pens running out. This raises the question whether the frequency of a particular action, or non-action, also influences our judgement of its causal contribution in a causal structure where the frequency itself does not make a difference to the outcome. The causal situation in the pen case is *cumulative*: The more often the action is performed, the more likely the final outcome is to happen. However, let’s imagine a case where at the end of each day the pen stock gets replenished to two pens. If only one of the agents repeatedly takes a pen on each day, then no problem occurs. But when both agents decide to take a pen on the same day, then there will be no pens left for emergency cases. This case represents a conjunctive causal structure – two actions are needed for the adverse outcome to occur. We were interested whether the frequency of an action still has an impact on our causal judgements in such a scenario:

1. Does the frequency of an action affect causal judgements in a conjunctive causal structure?

Norm-incorporating theories predict increased causal attribution to actions that violate moral norms, but make no prediction as to whether the frequency of the norm violation further accentuates its causal relevance. In terms of counterfactual dependence, if the joint actions of a frequent and a one-off moral norm violator lead to a negative outcome in a conjunctive causal structure, they are both equally causal – if either of them had not acted, the outcome would not have occurred (Halpern & Pearl, 2010). Furthermore, they are both also ‘equally abnormal’, because both of them violated the norm in the situation of the final outcome. Current norm incorporating theories leave open whether we mentally undo only the final vs. all previous norm-violating actions in order to assess the causal relevance of a norm-violating agent. This motivated our second research question:

2. Is a frequent vs. one-off norm violation assigned a different amount of causality in a conjunctive causal structure?

Gershman et al. (2016) show that if a certain action leads to a negative outcome in a new context, e.g. the door at the new office breaks when you turn the doorknob clockwise, this action will be judged less blameworthy when it has been

a habitual action before, e.g. when the door knob at home runs clockwise. In fact, we can think of situations in which an action that has been frequently done before is suddenly not allowed anymore, for example smoking in bars after the UK smoking ban in 2007. This raises the question whether a norm-violating act that has been permissively performed before is judged less causal than a repeated norm violation:

3. How does a frequent norm violation compare to an action that has been done frequently, but violated a norm only once?

The two experiments in this paper aim to address these questions.

Experiment 1

Experiment 1 examined the influence of the “frequency” and “morality” of an action on causal judgments in a conjunctive causal structure. Frequent and moral behavior were varied across the actions of one agent (‘Agent 1’) and held constant across actions of the other agent (‘Agent 2’). In order to manipulate frequent behavior realistically, we used the time frame of a week in which information about the action of the agents is successively presented day by day. Additionally, we varied whether the official norm indicating which agent is allowed to perform the respective action is either introduced right at the beginning of the scenario, or just before the final outcome (“time point of moral norm”). By varying the time point at which the moral norm was introduced we could test actions that either frequently violate a norm, or violate the norm for the first time but have been frequently performed before.

Methods

Participants 103 participants were recruited via Amazon Mechanical Turk (\$0.80 for ca. 15 min). Participants who answered more than four of the 32 manipulation check questions wrong were removed from the analysis, leaving 81 participants (male = 45, female = 36; mean age=34.78; SD=11.36, age range = 19-74).

Materials All conditions used variations of the same vignette. The scenario describes a week in a startup company, from Monday to Monday, with office days from Monday to Thursday. The company printer breaks as soon as it receives two printing orders at the same time [conjunctive causal structure with two actions], while individual printing does not make a causal difference. The participant acts as CEO with two co-workers, Agent 1 and Agent 2, whose names and gender are counterbalanced across scenarios.

Together with two former classmates from university, Agent 1 and Agent 2, you have founded a small startup company of which you are the CEO. You are in charge of the management and finances, while Agent 1 and Agent 2 manage the creative direction part. Your team works from Mondays to Thursdays. The printer in your office works fine, only when it receives two printing orders at the same time it crashes.

Agent 1 either uses the printer frequently during the week (“frequent behaviour”), or never prints during the week (“infrequent behaviour”). Furthermore, there is a company policy that determines that Agent 1 and Agent 2 are either allowed to use the printer anytime they want, or only on selected days.

As the CEO of your startup, you officially rule that both Agent 1 and Agent 2 are allowed to use the printer whenever they want [Agent 1 is only allowed to print on Tuesdays and Thursdays, and Agent 2 only on Mondays and Wednesdays.]

By printing on e.g. a Monday, Agent 1 conforms to the company policy in case of the liberal printer policy (“moral norm conformation”) or violates the company policy in case of the selective policy (“moral norm violation”). Combined with the frequency manipulation, Agent 1 therefore conforms to or violates the policy, and does so frequently or infrequently. In contrast, Agent 2 prints frequently during the week, and always only on the days on which she is allowed to print. Each day of the week is successively presented.

On Monday, X uses the printer for printing the new flyer design.
 On Tuesday, X uses the printer for printing the program outline.
 On Wednesday, X uses the printer for the printing of the new logo.
 On Thursday, X uses the printer for printing out the schedule.

After the presentation of one week, i.e. the Monday after the start of the scenario, both agents simultaneously send printing orders to the printer and it crashes.

On the following Monday, you come into the office and the printer is broken. Both Agent 1 and Agent 2 have sent printing orders to the office printer today. This is bad because your start up currently does not have the budget to afford a new high quality colour printer.

The company policy is either introduced at the beginning of the scenario, as shown in the previous examples, or on Thursday, i.e. just before the final outcome (“time point of moral norm”)

Design and Procedure The experiment was designed as 2 (Frequency) × 2 (Moral Norm) × 2 (Time point of moral norm) within subject paradigm. Participants saw all eight variations of the scenario in a randomized order, and each office day was presented successively on a single slide. The participants then had to judge the causal contribution of both agents to the outcome on a 10-point causal rating scale (“To what extent did Agent X cause the outcome?”: 0”- ‘None at all’; “10”- ‘Fully’). Recent studies have highlighted that the term ‘cause’ in the test questions is ambiguous and can refer to both the causal mechanism and the agent’s accountability (Samland et al., 2015). Therefore, we added a 7-point counterfactual relevance agreement scale (“If Agent X had not printed, the problem would not have occurred.”; “1” - ‘Strongly disagree’; “7” - ‘Strongly agree’) to directly test counterfactual reasoning. In addition, four manipulations check questions about the moral norm (“In the scenario you have just read, was Agent X allowed to print on Monday?”) and the frequency of the agents’ actions (“In the scenario you

have just read, did Agent X typically use the printer?”) with the answer options (‘yes’, ‘no’) were given.




Figure 1. Experimental design with the norm introduced on beginning of the scenario (above) and on Thursday, i.e. at the end of the scenario (below).

Results

Causal Rating | Agent 1 Causal ratings for Agent 1 were higher when they violated the moral norm (M=6.35; SD=2.43) than when they conformed to it (M=4.53; SD=1.60), $F(1, 80) = 67.71, p < .001, \eta_p^2 = .46$. In contrast, when Agent 1 acted only once, they were judged as less causal (M=4.80, SD=2.46) than when they performed a frequent behaviour (M=6.06, SD=1.71), $F(1, 80) = 25.08, p < .001, \eta_p^2 = .24$. The time point of the norm did not reveal any significant effects ($p=.83$).

Causal Rating | Agent 2 Agent 2 was seen as less causal when Agent 1 violated a moral norm (M=3.61, SD=2.33) compared to when Agent 1 conformed to the norm (M=5.40, SD=1.78), $F(1, 80) = 53.33, p < .001, \eta_p^2 = .40$. Likewise, Agent 2 was seen as more causal when Agent 1 printed one-off (M=4.89, SD=2.32), compared to when Agent 1 printed frequently (M=4.13, SD=1.60), $F(1, 80) = 13.46, p < .001, \eta_p^2 = .14$. The time point of the norm did not reveal any significant effects ($p=.16$).

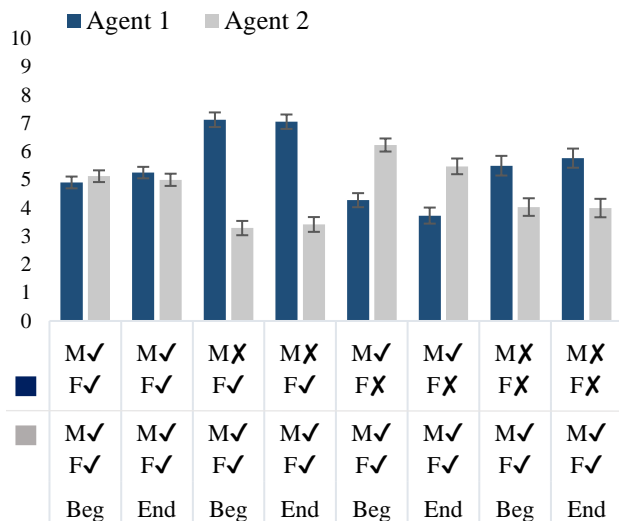


Figure 2. Causal Rating. “M” indicates the *moral norm* and “F” the *frequency* of behaviour, with ‘✓’ for ‘conforming/frequent’ and ‘✗’ for ‘violating/one-off’. The error bars represent ± 1 SE of the mean.

Counterfactual Relevance | Agent 1 The agreement ratings for the counterfactual relevance of Agent 1 were affected by both Agent 1’s moral behaviour, $F(1, 80) = 21.60, p < .001, \eta_p^2 = .21$, and frequency of action, $F(1, 80) = 17.86, p < .001, \eta_p^2 = .18$. The agreement for Agent 1 as relevant counterfactual was higher when they violated a moral norm ($M=5.20, SD=1.47$) than when they did not ($M=4.75, SD=1.44$) and lower when they acted one-off ($M=4.69, SE=1.63$) compared to when they acted frequently ($M=5.26, SD=1.32$).

Counterfactual Relevance | Agent 2 The violation of morality by Agent 1 decreased agreement with Agent 2 as counterfactually relevant ($MD=-.45; SD=0.13$), $F(1, 80) = 12.67, p < .001, \eta_p^2 = .14$. Agent 2 was also seen less counterfactually relevant when Agent 1 acted frequently ($MD=-.32; SD=0.13$), $F(1, 80) = 12.67, p=.006, \eta_p^2 = .14$

Moral Norm Violation A 2 (Frequency) \times 2 (Time point of moral norm) ANOVA for norm violating actions revealed a significant difference for the frequent norm violation ($M=7.09, SD=1.24$), vs. infrequent norm violation ($M=5.62, SD=1.36$) of Agent 1 on causal judgements, $F(1, 80) = 24.60, p < .001, \eta_p^2 = .24$. A corresponding reverse rating was found for Agent 2 when Agent 1 violated the moral norm frequently ($M=3.24, SD=1.25$) vs. once ($M=3.99, SD=1.33$), $F(1, 80) = 8.43, p = .005, \eta_p^2 = .09$. Counterfactual relevance ratings for Agent 1 were higher when Agent 1 frequently violated the norm ($M=5.49, SD=.15$) vs. once once ($M=4.92, SD=.20$), $F(1, 80) = 16.17, p < .001, \eta_p^2 = .17$. Agent 2 was seen less counterfactually relevant when Agent 1 frequently violates the norm, compared to violating the norm once ($MD=-.47, SD=.18$), $F(1, 80) = 7.23, p = .009, \eta_p^2 = .08$.




Figure 3. Counterfactual relevance rating. “M” indicates the *moral norm* and “F” the *frequency* of behaviour, with ‘✓’ for ‘conforming/frequent’ and ‘✗’ for ‘violating/one-off’. The error bars represent ± 1 SE of the mean.

Discussion

Experiment 1 confirmed the effect of moral norms on causal attribution (Knobe & Fraser, 2008), showing that an agent who violates a moral norm is judged as more causal and counterfactually relevant than a norm-conforming agent. We also found a reversed norm influence for agent-level statistical norms, with frequent behaviour being judged more instead of less causal than one-off actions. We found the same effects for counterfactual relevance ratings, showing that moral norm violations and frequent behaviour increase the extent to which an action is seen as counterfactually relevant to the outcome. An increase in causal attribution to Agent 1 due to norm violating and/or frequent behaviour was accompanied by a decrease in causal attribution to Agent 2, known as “causal superseding” (Kominsky et al., 2015). We did not find an interaction between the immoral norm behaviour and the time point at which the norm was introduced in the scenario. Hence, an action that has been violating the norm from the beginning is judged as causal as an action that has been frequently performed before, but only violated a norm after the new introduction of a moral norm.

Our finding that frequent behaviour increased causal attribution is consistent with Sytsma et al. (2012), but was surprising insofar as the causal structure of our scenario was designed to be causally insensitive to an individual action that occurs frequently, i.e. using the printer often. However, it might be argued that the influence of frequent vs. one-off actions in a conjunctive causal structure depends on the knowledge that both agents have about each other’s behaviour. For example, when the agent who usually never uses the printer suddenly prints in a conjunctive causal structure, she might be seen as more causally responsible for the outcome if she is aware that the other agent frequently

prints. In addition, despite our attempt to implement a conjunctive causal structure, participants might have interpreted frequent behaviour in our scenario as gradually damaging the printer. In order to further investigate the role of knowledge, and accentuate the difference between a cumulative and conjunctive causal structure, we conducted a second experiment.

Experiment 2

Experiment 2 examined the influence of frequent behavior in dependence of the underlying causal structure and state of knowledge of the agents.

Methods

Participants 102 participants were recruited via Amazon Mechanical Turk (\$0.80 for ca. 15 min), and after removing participants who had more than four of the 32 check questions wrong, the data of 64 participants were analyzed (male = 34, female = 30; mean age=33.56; SD=11.92, age range = 21-74).

Materials The same design as in Experiment 1 was used, except that the moral norm manipulation was removed by setting the moral status of all actions to “moral norm conforming” by a company policy that allowed the two agents to use the printer anytime. The frequency of the actions of Agent 1 was varied while holding the frequency of Agent 2 constant. Hence, either both of them frequently used the printer on different days during the week, or only Agent 2 used the printer. The underlying causal structure of the scenario was also varied by either containing a conjunctive causal structure (Experiment 1), or a cumulative causal structure, i.e. every action gradually increases the likelihood for the outcome to happen

The printer is quite worn out, and every printing order sent strains the printer system a bit more.

Finally, we varied whether the agents know about each other’s behavior.

The printer management system does [does not] display the current user balance so that Agent 1 and Agent 2 know [do not know] who prints how much each day.

Design and Procedure The experiment was designed as a 2 (Frequency) × 2 (Causal Structure) × 2 (Knowledge) within Subject Design. The participants saw each of the eight scenarios in a randomized order, and answered two causal strength and counterfactual relevance questions, using the same scales as in Experiment 1. Following this, they were asked two manipulation check questions about the typical behavior (“In the scenario you have just read, did Agent X typically use the printer?”) a knowledge manipulation question (“In the scenario you have just read, does the printer management system display the user balance?”), and a causal structure check question (“In the scenario you have just read, under which conditions does the printer crash?”), with the

options “Two or more printing orders at the same time” or “Overuse”.

Results

Causal Rating | Agent 1 Causal ratings for Agent 1 were generally higher when they acted frequently (M=5.31; SD=1.30) than when they acted one-off (M=2.90; SD=1.89), $F(1, 63) = 103.78, p < .001, \eta_p^2 = .62$. However, the increase in causal contribution when Agent 1 used the printer frequently was greater in the cumulative causal structure than in the conjunctive causal structure, $F(1, 63) = 103.78, p = .001, \eta_p^2 = .16$. No significant effect for knowledge was found ($p = .48$).

Causal Rating | Agent 2 Agent 2 was seen as less causal when Agent 1 acted frequently, (M=5.32, SD=1.30), but more causal when Agent 1 performed a one-off action (M=6.93, SD=1.90), $F(1, 63) = 69.15, p < .001, \eta_p^2 = .52$. The increase in causal attribution to Agent 2 when Agent 1 did not act frequently was greater in the cumulative than in the conjunctive causal structure, $F(1, 63) = 39.64, p < .001, \eta_p^2 = .37$. There was no significant effect for knowledge ($p = .90$).

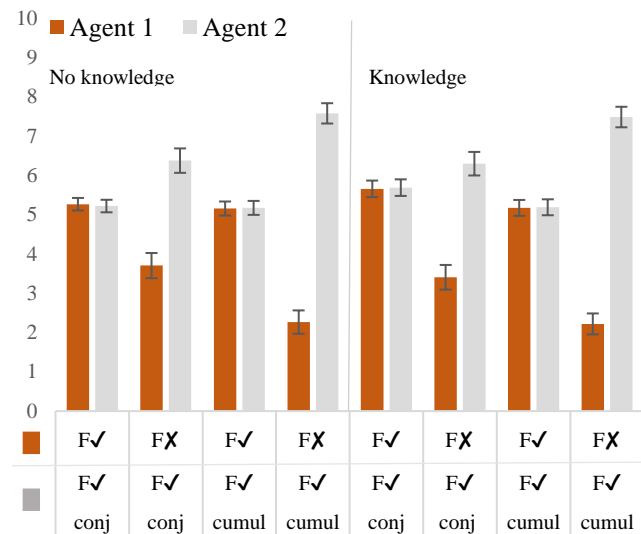


Figure 4. Causal Rating. “F” indicates the frequency of behaviour, with ‘1’ for ‘frequent’ and ‘0’ for ‘one-off’. The error bars represent ± 1 SE of the mean.

Counterfactual Relevance Agent 1 was seen as more counterfactually relevant when they acted frequently (M=4.86; SD=.1.33) versus one-off (M=3.93; SE=1.56), $F(1, 63) = 29.87, p < .001, \eta_p^2 = .32$, and the increase in counterfactual relevance when acting frequently was greater in the cumulative structure, $F(1, 63) = 6.90, p = .011, \eta_p^2 = .10$. Agent 2 is assigned less counterfactual relevance when Agent 1 acts frequently (MD=-.56; SD=.197, $F(1, 63) = 12.19, p = .02, \eta_p^2 = .16$. There was no significant effect for the knowledge factor ($p = .50$).




Figure 5. Counterfactual Relevance Rating. “F” indicates the frequency of behaviour, with ‘1’ for ‘frequent’ and ‘0’ for ‘one-off’. The error bars represent ± 1 SE of the mean.

Discussion

Experiment 2 confirmed the effect of frequent behavior on causal judgments. Frequent behavior was seen as more causal than one-off behavior, independent of the underlying causal structure. However, participants were aware that the change from frequent behaviour to a one-off action led to a greater reduction in causal contribution in the cumulative causal structure than in a conjunctive structure. The effect of frequent behaviour was independent of whether or not the agents knew about each other’s behavior.

General Discussion

Our two experiments found evidence for an increase in the causal attribution to an agent who acts in line with their frequent behaviour, even when the frequency of this behavior does not make a causal difference to the outcome. Our results present a challenge to the norm account of causal judgment, showing that a one-off action that deviates from frequent behavior is judged less instead of more causal for a negative outcome. Since the knowledge about the other agent’s behavior did not make a difference to causal judgements, despite being crucial for estimating the outcome in a conjunctive causal structure, we rule out that people might have inferred bad intentions or foreseeability of the outcome from typical behavior (Lagnado & Channon, 2008). Hence, we believe that the fact that people assign increased causality to a frequent, typical action goes against predictions of norm theories and shows the need for a new account of frequency of actions in causal attribution. Current counterfactual and structural equation models (Halpern & Pearl, 2005; Chockler & Halpern, 2014) fail to account for the asymmetry of causal attribution between a frequent vs. one-off actor in a conjunctive causal structure, given that previous behavior does not change the counterfactual dependency of the two actions in the actual situation.

We draw two conclusions from this. First, we argue for the need to include the previous history of actions of an agent

into counterfactual accounts of causal judgments. Second, we argue that an extension of current counterfactual theories is needed in order to capture the influence of frequency. One such extension could be probabilistic (Fenton-Glynn, 2016). While frequent behavior does not differ from one-off behavior in terms of single counterfactual dependencies, it does so probabilistically. If, in a counterfactual world, the other agent acts at a random time point in a conjunctive structure, the outcome is more likely to occur if I frequently perform the other action needed for the outcome. In contrast, it is less likely to occur if I act infrequently. The raised probability of the outcome due to my frequent behavior can even increase when we also vary whether the other agent acts frequently vs. one-off. As a result, we argue that in addition to counterfactually testing whether the undoing of an action makes a difference, we also need to test whether the variation of the frequency of the action would make difference to the likelihood of the outcome to occur.

Our results show that a frequent norm violation is judged more causal than a one-off moral norm violation. This is despite the fact that in the actual situation of the outcome, they are both equally counterfactually relevant, as well as equally. Hence, we think that current norm theories need to include the frequency of previous norm violations in order to fully capture the influence of moral norms on causal attributions. However, the time point at which the moral norm is introduced into the scenario does not change the way we assign causality to actions that violate this norm. To conclude, the influence of the frequency of an action proves to be a crucial factor in the attribution of causality, and calls for new theoretical frameworks of causation. We argue that this framework has to be probabilistic.

References

- Alicke, M. D., Rose, D., & Bloom, D. (2011). Causation, Norm Violation, and Culpable Control. *Journal of Philosophy*, 108(12), 670–696.
- Fenton-Glynn, L. (2016). A Proposed Probabilistic Extension of the Halpern and Pearl Definition of “Actual Cause.” *The British Journal for the Philosophy of Science*, 78, 85–109.
- Gershman, S. J., Gerstenberg, T., Baker, C. L., & Cushman, F. A. (2016). Plans, Habits, and Theory of Mind. *PLoS One*, 11(9), e0162246.
- Gerstenberg, T., & Lagnado, D. A. (2010). Spreading the blame: The allocation of responsibility amongst multiple agents. *Cognition*, 115(1), 166–171.
- Hart, H. L. A., & Honoré, T. (1985). *Causation in the Law*, 2nd ed., Oxford: Clarendon.
- 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. (2005). Causes and Explanations: A Structural-Model Approach. Part I: Causes. *The British Journal for the Philosophy of Science*, 56(4), 843–887.
- Hitchcock, C., & Knobe, J. (2009). Cause and Norm. *Journal of Philosophy* 106 (11): 587–612.
- Icard, T., Kominsky, J., and Knobe, J. (2016). Causality, Normality, and Sampling Propensity. In : *Proceedings of the 38th Annual Meeting of the Cognitive Science Society*.
- 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. *Moral Psychology*. In: Walter Sinnott-Armstrong (ed.), *Moral Psychology*. MIT Press
- Kominsky, J. F., Phillips, J., Gerstenberg, T., Lagnado, D., & Knobe, J. (2015). Causal superseding. *Cognition*, 137, 196–209.
- Lagnado, D. A., & Channon, S. (2008). Judgments of cause and blame: The effects of intentionality and foreseeability. *Cognition*, 108(3), 754–770.
- Roxborough, C., & Cumby, J. (2009). Folk psychological concepts: Causation I. *Philosophical Psychology*, 22(2), 205–213.
- Samland, J., & Waldmann, M. R. (2015). Highlighting the causal meaning of causal test questions in contexts of norm violations. In D. C. Noelle et al. (Eds.), *Proceedings of the 37th Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.
- 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), 814–20.