



Graded causation and moral responsibility

Vera Hoffmann-Kolss¹ · Matthias Rolffs¹

Received: 4 January 2023 / Accepted: 20 February 2024
© The Author(s) 2024

Abstract

Theories of graded causation attract growing attention in the philosophical debate on causation. An important field of application is the controversial relationship between causation and moral responsibility. However, it is still unclear how exactly the notion of graded causation should be understood in the context of moral responsibility. One question is whether we should endorse a proportionality principle, according to which the degree of an agent's moral responsibility is proportionate to their degree of causal contribution. A second question is whether a theory of graded causation should measure closeness to necessity or closeness to sufficiency. In this paper, we argue that we should indeed endorse a proportionality principle and that this principle supports a notion of graded causation relying on closeness to sufficiency rather than closeness to necessity. Furthermore, we argue that this insight helps to provide a plausible analysis of the so-called 'Moral Difference Puzzle' recently described by Bernstein.

1 Introduction

When we seek explanations, be it in everyday situations or scientific contexts, we typically try to identify the causes of the phenomena which are of interest to us. In most cases, a single phenomenon will have several causes, and we are often not only interested in identifying them but also in assessing how important they were for the occurrence of the effect. For instance, an intensely debated topic in current climate science is the question of how significant the causal contribution of certain types of action, such as mobility or production of food, is to global warming. Similarly, at the

✉ Vera Hoffmann-Kolss
verahoffmann@gmx.de

Matthias Rolffs
matthias.rolffs@unibe.ch

¹ University of Bern, Institute of Philosophy, Laenggassstrasse 49a, Bern 3012, Switzerland

level of individuals, medical scientists are often interested in determining the degree to which smoking, dietary habits and environmental toxins contribute to a person developing cancer. In other words, everyday accounts, as well as scientific contexts, strongly suggest that causation allows for degrees: if events c_1 and c_2 both cause a further event e , it may well be the case that c_1 is more of a cause of e than c_2 is.

It is, therefore, not surprising that theories of graded causation attract growing attention in the philosophical debate on causation and that philosophers seek to determine the conditions of adequacy that an account of degrees of causation is supposed to fulfil. With respect to non-graded causation, it is widely agreed that a unified concept that describes either some fundamental relation in the world or at least some fundamental concept of human thinking is available (prominent exceptions to this view are Godfrey-Smith, 2009; Hall, 2004; Psillos, 2007). Concerning the notion of graded causation, by contrast, many philosophers do not even assume that there is a unified concept. As Kaiserman puts it:

...in describing A as ‘more of a cause’ of some effect than B, we might be saying that A made more of a difference (either to the effect or to its probability), that it came closer to making a difference, that it contributed to more causings of the effect, that it contributed to these causings to a larger degree, or indeed something else entirely. Considered apart from their potential applications, none of these measures is any ‘better’ than the others – they are simply measuring different things. (Kaiserman, 2018: 5)

In addition to the different notions of causal contribution mentioned by Kaiserman, there seems to be a fundamental difference between theories of causal contribution that aim to satisfy what Sartorio has dubbed the ‘Necessity Criterion’ and theories supposed to fulfil what she calls the ‘Sufficiency Criterion’:

Necessity Criterion: How much a cause contributes to an effect is a matter of how close it comes to providing a *necessary* condition for an effect.

Sufficiency Criterion: How much a cause contributes to an effect is a matter of how close it comes to providing a *sufficient* condition for an effect. (Sartorio, 2020: 349–350)

An account of graded causation can typically meet either the Necessity Criterion or the Sufficiency Criterion, but not both. This raises the question of which of these two criteria should be imposed as an adequacy condition on a theory of graded causation. We assume that it depends on the function the notion of causal contribution is intended to fulfil. What we are concerned with is, accordingly, not a descriptive conceptual analysis of the notion of graded causation. We are not primarily interested in how people actually use locutions like ‘A is more of a cause than B’. Instead, the present project can be seen as an exercise in conceptual engineering (Brun, 2016).

We probe different concepts of graded causation with respect to their suitability to serve a certain function.¹

The function of interest to us concerns the relationship between causation and responsibility. This is one of the most important applications of the notion of graded causation. We try to determine the contribution of certain actions to climate change not only for purely academic reasons but also because we want to hold actors, such as companies, responsible for damages resulting from climate change (thus trying to force them to reduce their CO₂ emissions). For example, when several companies are sued for jointly causing some environmental damage by emitting effluents into a river, the degree to which each should be held responsible for the damage seems to crucially depend on the degree to which they causally contributed to it (Braham & Hees, 2009: 324–5). In general, if several agents contribute to some (harmful) outcome *e*, their contributions may differ with respect to how much they were causally relevant to *e*, and this partially determines the extent to which they are held responsible or liable for *e* (see, e.g., Hart & Honoré, 1959; Kaiserman, 2017; Moore, 2009: 3–19, 118–121; Tadros, 2018).

The consideration that responsibility depends, at least to a certain extent, on causal contribution leads to a plausible principle that Bernstein calls ‘Proportionality’:

Proportionality: The degree of responsibility of an agent *A* for an outcome *o* is (*ceteris paribus*) proportionate to the degree to which *A*’s actions (or omissions) caused *o* (Bernstein, 2017).

Bernstein argues that accepting Proportionality leads to significant confusion concerning moral responsibility, namely the so-called ‘Moral Difference Puzzle’ (Bernstein, 2017; see also Sartorio, 2020). One central aim of this paper is to defend Proportionality against this charge of creating confusion.²

More specifically, we aim to establish the following three claims: first, Proportionality, combined with more general considerations about the function of the notion of graded causation, leads to two adequacy conditions that a theory of graded causation must fulfil to account for the relationship between degrees of causation and degrees

¹ There are several X-Phi studies that seem to presuppose that people actually use a graded notion of causation (e.g., Gill et al., 2022; Quillien & Barlev, 2022). These studies specifically examine the role of norms for causal judgement. However, these studies are not directly relevant for present purposes, mainly for two reasons: first, as they examine the actual usage of ordinary speakers, they cannot tell us which concept of graded causation people *should* use. Second, the questionnaires in these studies typically do not directly ask participants about graded causal statements (such as ‘*A* is more of a cause than *B*’). Rather, they ask for degree of agreement with non-graded causal statements (such as ‘*A* caused *B*’). Therefore, it is problematic to interpret their results as directly relevant to theories of graded causation.

² A more general objection to Proportionality is that concepts of graded causation which are motivated by specific philosophical accounts of causation, such as Dowe’s (2000) conserved quantity account or a standard counterfactual account à la Lewis (1973), do not deliver the correct verdicts when inserted into Proportionality (Bernstein, 2017; Demirtas, 2022b). We agree that some independently motivated accounts of (graded) causation may lead to implausible verdicts about moral responsibility when combined with Proportionality. However, this is compatible with the argument of this paper. In line with the conceptual engineering approach mentioned above, we develop constraints on an adequate concept of graded causation (and these constraints may or may not be compatible with the consequences of specific accounts of graded causation).

of responsibility. Second, these adequacy conditions argue for accounts based on the Sufficiency Criterion rather than the Necessity Criterion. Third, given this insight, Proportionality does not create confusion about moral responsibility but allows for a plausible analysis of Bernstein's (2017) Moral Difference Puzzle.

In Sect. 2, we formulate the two above-mentioned adequacy conditions for a theory of graded causation and argue that these conditions support the Sufficiency Criterion rather than the Necessity Criterion. We then argue, in Sect. 3, that the acceptance of the Sufficiency Criterion leads to a version of Proportionality that does not create confusion about moral responsibility. We conclude with a summary and a brief comment on the implications of our argument for the debate on moral luck.

2 Proportionality and redundant causation

The current literature on graded causation includes both quantitative and comparative accounts. A quantitative measure of degrees of causation is a function that assigns a concrete numerical value to each cause in a given causal structure. A comparative measure of degrees of causation ranks causes of a given effect according to their relative contribution to it. Since most accounts of degrees of causation currently on the market are quantitative (e.g., Braham & Hees, 2009; Chockler & Halpern, 2004; Ferey & Dehez, 2016; Northcott, 2013; Sprenger, 2018), we focus on this type of account in what follows.³⁴

Degrees of causation are of interest when considering causal scenarios where one effect has several causes. Here, a fundamental distinction to be drawn is between non-redundant and redundant structures. Suppose that an effect e has two causes c_1 and c_2 . Then the causal structure including c_1 , c_2 , and e is redundant iff if c_1 had occurred without c_2 , c_1 would still have caused e , so that e would still have occurred, and the analogous condition holds if c_2 had occurred without c_1 (Schaffer, 2003: 23, see also Lewis, 1986: 193). Otherwise, the structure is non-redundant. If e has more than two causes c_1, \dots, c_n , the causal structure including c_1, \dots, c_n , and e is redundant iff if either of the causes c_1, \dots, c_n had not occurred, but the others had, then the others would still have caused e , so that e would still have occurred. Otherwise, the struc-

³ It should be noted, however, that one of the most prominent contributions to the debate on graded causation, the account proposed by Halpern and Hitchcock (2015), is comparative.

⁴ A further conceptual distinction is between *causing outcome o to degree x* and *causing some portion of outcome o* . Demirtas discusses the following example: Suzy and Billy pick a total of 1,000 apples together in an orchard, but whereas Suzy picks 600 apples, Billy picks only 400 apples. Here, it seems natural to say that Suzy caused 60% of the outcome, and Billy caused 40% of the outcome. In this case, it also seems natural to translate this into degrees of causation: Suzy contributed to the outcome to degree 0.6, whereas Billy contributed to degree 0.4 (Demirtas, 2022b). However, that does not mean that the notion of *causing outcome o to degree x* can be equated with the notion of *causing some portion of outcome o* . Many effects are not quantifiable and distributable. If an outcome is not distributable, it usually makes no sense to say that several actions caused some portion of it. However, it can make perfect sense to say that an agent contributed to it to degree x (for further discussion see Demirtas (2022b); Kaiserman, (2017: 4) introduces a similar conceptual distinction). In what follows, we only focus on the notion of *causing outcome o to degree x* .

ture is non-redundant.⁵ A paradigmatic example of a causally redundant structure is a classical overdetermination scenario involving two or more sufficient causes of the same effect.

In what follows, we examine the consequences of Proportionality (restated here for convenience) for several cases of non-redundant as well as redundant causation.

Proportionality: The degree of responsibility of an agent A for an outcome o is (*ceteris paribus*) proportionate to the degree to which A's actions (or omissions) caused o (Bernstein, 2017).

Note that Proportionality contains a *ceteris paribus* clause. Of course, there is more to moral responsibility for an outcome than the degree of causal contribution: one can contribute to an outcome to a very high degree but still not be responsible for it. If certain other conditions for moral responsibility – for example, conditions concerning the epistemic situation of the agent or the availability of reasonable alternatives to the action – are not fulfilled, the agent will not be responsible at all (see Kaiserman (2021: 3602–3603) for discussion). In our examples, we assume that all other conditions for moral responsibility are fulfilled. Insofar as these other conditions allow for degrees, we assume that they are all fulfilled to the same degree.

Let us now look at a simple, non-redundant scenario:

Pollution-1 (joint causation – different amounts): Two companies, A and B, dump chemicals into a river, and as a result, a fish species inhabiting the river dies out. In the given circumstances, ten units of chemicals are necessary and sufficient for the extinction of the fish species. Company A dumps nine units; company B dumps only one unit.

Pollution-1 provides strong intuitive support for Proportionality. In Pollution-1, it seems plausible that company A is more responsible for the extinction of the fish species than company B. This difference in moral responsibility can be explained by a difference in causal contribution: company A is more responsible because it contributed to the harm to a higher degree. Proportionality can explain this intuition, whereas if Proportionality were rejected, there would be no explanation for the difference in responsibility of the two companies – at least, provided that all other circumstances and background conditions are held equal.

Next, consider two further non-redundant scenarios:

Pollution-2 (necessary and sufficient causation): A single company, A, dumps ten units of chemicals into a river, and as a result, a fish species inhabiting the river dies out. In the given circumstances, the amount of chemicals emitted by company A is both necessary and sufficient for the extinction of the fish species.

⁵ There may be mixed cases, where some of the causes of e, c_1, \dots, c_n , are redundant while others are not. However, since none of the scenarios discussed in this paper fall into this category, we will leave this complication aside.

Pollution-3 (joint causation – same amount): Two companies, A and B, dump chemicals into a river, and as a result, a fish species inhabiting the river dies out. In the given circumstances, ten units of chemicals are necessary and sufficient for the extinction of the fish species. Company A dumps five units, and company B also dumps five units.

Plausibly, company A's responsibility in Pollution-3 should be lower than in Pollution-2. After all, company A's action in Pollution-2 is sufficient on its own for killing the fish species, whereas in Pollution-3, company A's action on its own would not have the same damaging effect. If company A in Pollution-3 had acted in isolation, the fish species would not have died. It, therefore, does not seem appropriate to assign them the same degree of responsibility as in Pollution-2.⁶ These observations motivate the following condition of adequacy for theories of graded causation:

(a) An account of graded causation should not imply that all causes necessary for the occurrence of a given effect are automatically assigned the same value.

In Pollution 1, company A's action and company B's action are both necessary for the extinction of the fish species. If company A had not dumped nine units, the fish species would not have gone extinct. And if company B had not dumped 1 unit, the fish species would not have gone extinct either. Therefore, company A's action and company B's action are equally close to being necessary for the harm, as they are both strictly necessary. Any account of graded causation that automatically ascribes maximal contribution to necessary causes cannot differentiate between company A's contribution and company B's contribution in Pollution-1. However, it is highly plausible that company A is responsible for the extinction of the fish species to a higher degree than company B and that this difference in responsibility is due to a difference in causal contribution.

As for the comparison between Pollution-2 and Pollution-3, a parallel problem arises. If all necessary causes are assigned the same degree of causal contribution, company A in Pollution-3 contributes just as much as in Pollution-2; in both cases, the contribution is strictly necessary. However, in the context of responsibility ascriptions, we should be able to differentiate between company A in Pollution-2 and company A in Pollution-3: in Pollution-2, company A is more responsible for the harm than in Pollution-3.

It should be noted that condition (a) is, of course, compatible with scenarios where an effect has several causes that are all necessary for its occurrence and that are all assigned the same value, that is, that all make the same causal contribution. It is plausible to assume, for instance, that in Pollution-3, the actions of company A and company B, who both dump five units of chemicals, should be assigned the same value. The point of condition (a) is just that such an assignment of values should not be *automatically implied* by the fact that the actions are both necessary.

Let us now turn to cases of redundant causation. Compare the following two cases:

⁶ This is in line with Kaiserman's (2021: 3605) account of a similar pair of cases.

Pollution-4 (overdetermination – two causes): Two companies, A and B, each dump ten units of chemicals into a river, and as a result, a fish species inhabiting the river dies out. In the given circumstances, the amount of chemicals emitted by either company would have been sufficient to kill the fish species.

Pollution-5 (overdetermination – three causes): Three companies, A, B, and C, each dump ten units of chemicals into a river, and as a result, a fish species inhabiting the river dies out. In the given circumstances, the amount of chemicals emitted by any of the companies would have been sufficient to kill the fish species.

If we compare Pollution-4 to Pollution-5, company A should plausibly have the same degree of responsibility in both cases. After all, company A's action is (intrinsically) the same, and responsibility should not be reduced by the mere presence of yet another redundant cause. Otherwise, one might be in the grip of what Kaiserman dubs the 'pie fallacy': 'the fallacy of thinking that there is a fixed amount of responsibility for every outcome, to be distributed among all those, if any, who are responsible for it' (Kaiserman, 2021: 3597). From this perspective, company A in Pollution-5 is indeed less responsible than in Pollution-4 because we need to distribute the same amount of responsibility among three agents in Pollution-5 instead of two agents in Pollution-4.

However, we agree with Kaiserman that responsibility does not work this way: the mere fact that company C is also responsible for the death of the fish species does not force us to say that company A becomes less responsible. As Zimmerman notes: 'What an easy "out" that would be! Just invite a few more friends to participate' (Zimmerman, 1985: 119). Responsibility is not like a pie with a limited supply, and responsibility shared is not necessarily responsibility halved (Schaffer (2003: 30) shares this verdict). This observation motivates the following condition of adequacy for theories of graded causation:

(b) An account of graded causation should not imply that causal contribution is reduced if further independent causes are added to a scenario.

One might be skeptical about this condition if one shares Chockler and Halpern's (2004) intuition that degrees of causation are indeed diminished in structurally similar cases. Chockler and Halpern present the following voting example as a main motivation for their view of graded causation: '[S]uppose that Mr. B wins an election against Mr. G by a vote of 11–0. Each of the people who voted for Mr. B is a cause of him winning. However, it seems that their degree of responsibility should not be as great as in the case when Mr. B wins 6–5' (Chockler & Halpern, 2004: 93; note that Chockler and Halpern's use of 'degree of responsibility' corresponds to our use of 'degree of causation').

How should one deal with this apparent clash of intuitions? We think that there are two considerations that tip the balance in favor of the intuition supporting condition (b).

First, reconsider the comparison between Pollution-4 and Pollution-5. One might think that company A is somewhat excused for emitting chemicals into the river

because, given that company B and C also emitted chemicals, this behavior is less exceptional. This idea becomes even more salient when we consider situations in which more companies showcase the same kind of behavior. If, say, a thousand companies each dump ten units of chemicals into the river, each company's behavior might be seen as conforming to an established norm to the effect that dumping chemicals is allowed. Such normality considerations might also play a role in people's failure to reduce their greenhouse gas emissions, or to stop eating meat.⁷ Even if it is morally required to reduce one's personal greenhouse gas emissions, or to stop eating meat, people might be excused to some degree for failing to do so because emitting greenhouse gases and eating meat is the norm. It is difficult to be a saint in a world of sinners, and if it is especially difficult to do the right thing, this might somewhat excuse one's failure to do it (see Nelkin (2016) for an exploration of the connection between difficulty and moral responsibility). As we see it, such considerations of normality often go along with cases of massive overdetermination, but they are in no way essential to them. This provides an alternative explanation for the intuition expressed by Chockler and Halpern: if responsibility is diminished in such cases, this can be explained by normality considerations rather than by the causal structure of the case. But then, the intuition that responsibility is diminished in such cases is compatible with condition (b), since it does not commit one to the claim that the presence of further overdetermining causes in and of itself diminishes responsibility.

The second consideration in favor of condition (b) arises from the debate over non-graded causation. Schaffer considers classical overdetermination scenarios with two overdetermining causes and discusses whether, in such cases, each of the overdeterminers individually should count as a cause of the effect under consideration, or whether only the mereological sum of the overdeterminers, but not the overdeterminers individually, should count as a cause of the effect (Schaffer, 2003: 24). He clearly argues for the first option. One of his arguments is that overdeterminers have the same predictive and explanatory function as ordinary causes. In a classical overdetermination scenario, where two sufficient causes, c_1 and c_2 , cause an effect e , 'knowledge that c_1 occurs is sufficient to license a prediction to the occurrence of e ', and 'citing the occurrence of c_1 serves to explain the occurrence of e ' (Schaffer, 2003: 29).

These considerations are consistent with the conceptual engineering approach of this paper. A philosophical account of causation should take into account the function that the notion of causation is intended to serve. One of the central functions of causation is to figure in predictions and explanations. And although Schaffer's argument applies primarily to non-graded causation, it is plausible to apply it to graded causation as well: causes that have the same predictive and explanatory power should be assigned the same degree of causal contribution. This provides additional support for (b): if further independent causes are added to a structure, the predictive and explanatory power of the causes already included in the structure remains the same. But then their causal contribution is not reduced, and this is what condition (b) says.

Therefore, condition (b) is a reasonable condition of adequacy to impose on a theory of graded causation. First, it can be argued that there is an alternative expla-

⁷ Thanks to an anonymous reviewer for pointing out this analogy.

nation for intuitions to the contrary, as described by Chockler and Halpen. Second, we have shown that the condition is supported by considerations about the function that the concept of causation has in predictive and explanatory contexts, and these are independent of intuitions about the relation between causation and responsibility.

However, theories of graded causation that rely on the Necessity Criterion violate both conditions (a) and (b). According to the Necessity Criterion, how much a cause contributes to an effect is a matter of how close it comes to providing a necessary condition for this effect. Suppose (as is common) that causal contribution is measured on a scale ranging from 0 to 1. Then, if a cause is a necessary condition of an effect, it should be assigned the maximal value of 1, regardless of the size of its contribution. But then, all causes that are necessary for the occurrence of the effect under consideration will be assigned that maximal value, and condition (a) will be violated.

Compare now a non-redundant scenario in which all causes of an effect are necessary conditions for the occurrence of the effect to a redundant scenario in which the effect would still occur if some of the causes were eliminated. In the non-redundant scenario, the single causes are maximally close to being necessary conditions (since they *are* necessary conditions). In the redundant scenario, however, how close a cause is to being a necessary condition depends on how many other redundant causes are present: if redundant causes are added to a structure, the causes already contained in the structure become less close to being necessary conditions. However, this consequence violates adequacy condition (b).⁸

By contrast, an account of graded causation that is in accordance with the Sufficiency Criterion will meet conditions (a) and (b). In Pollution-1, where company A dumps nine units into the river, company B dumps one unit, and ten units are necessary and sufficient for the extinction of the fish species, company A's action is much closer to being sufficient for the extinction of the fish species than company B's action. If causal contribution depends on how close causes are to being sufficient, company A's action will not be ascribed the same value as company B's action – although they are both necessary for the occurrence of the effect. But then condition (a) is fulfilled: causes that are necessary for the occurrence of a particular effect are not automatically assigned the same value.

These considerations are compatible with the observation that causes are usually not sufficient conditions for their effects in a strict sense since causal chains can always be interrupted by disturbing factors – the environmental activist who manages

⁸ Not surprisingly, Chockler and Halpern's (2004) account is one of the most prominent accounts that rely on the Necessity Criterion. The basic idea underlying this account is that how close a cause comes to being a necessary condition for an effect depends on how many changes one would have to make to the causal structure to turn it into a necessary condition. For instance, in a redundant scenario with four redundant causes one would have to eliminate four causes in order to turn the remaining causes into necessary conditions. In a scenario with just one redundant cause (e.g., an overdetermination scenario with two sufficient causes), only one cause must be eliminated in order to turn the other one into a necessary condition. But then the causes in the latter scenario (involving just one redundant cause) are closer to being necessary conditions than the causes in the former scenario (involving four redundant causes). This consequence violates adequacy condition (b). And since Chockler and Halpern are committed to assigning the maximal value 1 to all causes in a structure that are necessary conditions for the occurrence of the effect, their account violates condition (a) as well. Braham and Hees (2009) propose an account that has an analogous consequence. Kaiserman (2016, 2018) develops an account that satisfies condition (a).

to pour an antidote into the river, the flood that dilutes the concentration of the poison or, more far-fetched, the good fairy that casts a beneficial spell to protect the fish. Therefore, the descriptions of the scenarios all contain the caveat ‘in the given circumstances’. This caveat is unproblematic because it is possible to compare several causes with respect to closeness to sufficiency, even if they are not sufficient for the occurrence of a certain effect on their own. One just has to keep the circumstances, that is, all other potentially relevant factors, fixed. Given a fixed set of background assumptions, action A can be closer to being sufficient for the occurrence of effect E than action B, even if neither A, B, nor their conjunction on their own would have been sufficient for the occurrence of E.

To see that an account of graded causation based on the Sufficiency Criterion also satisfies adequacy condition (b), note that how close a cause is to being sufficient for a certain effect in a given causal structure will not change when further independent causes are added to the structure. Formally, this pattern is familiar from classical logic as ‘weakening the antecedent’: if $(A \rightarrow B)$, then $((A \wedge C) \rightarrow B)$. However, while in classical logic, the pattern holds for all C, in the causal case, the condition that the additional causes are independent of the causes already included in the structure is crucial. Suppose, for instance, that company A dumps ten units of chemical waste and that this is sufficient to kill the fish species, as in Pollution-2. If further chemicals that do not causally interact with the chemicals dumped by company A are added, this will not weaken the causal impact of company A’s action. On the other hand, if chemicals are added that causally interact with the chemicals dumped by company A, the result might be different. For example, consider the case where company B adds five units of chemicals that poison the fish and reduce the effect of the chemicals dumped by company A so that the chemicals dumped by company A only have the effect of five units. Then, company A’s action is no longer sufficient to kill the fish, and the causal contribution of company A is reduced. However, this case violates the condition that further *independent* causes are added since there is a causal interaction between company A’s and company B’s actions. If the condition that the causes added must not interact with the causes already included in the scenario is satisfied, weakening of the antecedent also holds in the causal case, and an account of graded causation based on the Sufficiency Criterion satisfies condition (b).

Therefore, if Proportionality is accepted, an account of graded causation should be in accordance with the Sufficiency Criterion rather than the Necessity Criterion. In the next section, we argue that this insight can help to shed light on a puzzle discussed in the literature.

3 A puzzle case for proportionality

Compare the following pair of cases, due to Bernstein:

Victim: Two independently employed assassins, unaware of each other, are dispatched to eliminate Victim. Being struck by one bullet is sufficient to kill Victim. Each assassin shoots, and Victim dies.

Hardy Victim: Two independently employed assassins, each unaware of the other, are dispatched to eliminate Victim. Unbeknownst to both assassins, Victim is particularly hardy and requires two bullets for his demise. Each assassin shoots, and Victim dies. (Bernstein, 2017: 165; see also Sartorio, 2020: 352 for a similar pair of cases)

Bernstein takes these cases to reveal ‘a puzzle about the relationship between degrees of causation and degrees of responsibility’ (Bernstein, 2017: 165) and ‘indeterminacy in the causal relations deployed in assessments of moral responsibility’ (Bernstein, 2017: 166). The reason is that, according to Bernstein, there are two contradictory intuitions about the two cases that both seem plausible. On the one hand, one might think that the assassins in Victim causally contribute more than the assassins in Hardy Victim since their actions are sufficient. On the other hand, one might think that the assassins in Hardy Victim contribute more than the assassins in Victim since their actions are necessary. In other words, ‘there are multiple causal relations that can be employed in moral assessments of differential causal contributions, and there are no clear, principled rules for which type of causal relation should be used’ (Bernstein, 2017: 172).

In what follows, we argue that this puzzle can be resolved. To see this, note, first, that two important questions about the comparison between these cases should be distinguished. Each question allows for three different answers.

First, the moral question: how does the assassin’s moral responsibility in Victim compare to their moral responsibility in Hardy Victim? Bernstein (2017: 165) dubs this question ‘the Moral Difference Puzzle’. Here, the three options are:

No moral difference. Each assassin is equally morally responsible for the victim’s death in Victim and in Hardy Victim.

Moral difference: greater. Each assassin in Victim is more morally responsible for the victim’s death than each assassin in Hardy Victim.

Moral difference: lesser. Each assassin in Victim is less morally responsible for the victim’s death than each assassin in Hardy Victim. (Bernstein, 2017: 166)

Second, the causal question: how do the causal contributions of each assassin’s action in Victim compare to the causal contributions of each assassin’s action in Hardy Victim? Here, the three options are:

No causal difference. Each assassin causally contributes the same to the victim’s death in Victim and in Hardy Victim.

Causal difference: greater. Each assassin in Victim causally contributes more to the victim’s death than each assassin in Hardy Victim.

Causal difference: lesser. Each assassin in Victim causally contributes less to the victim's death than each assassin in Hardy Victim.

Prima facie, Proportionality implies that the answer to the moral question must line up with the answer to the causal question. For example, if Proportionality is true, it would be inconsistent to opt for Causal difference: lesser and Moral difference: greater.

We take it that there is no independently obvious answer to the moral question. Rather, even before considering the causal question, the moral comparison between the two cases is genuinely puzzling (Bernstein (2017: 165) and Sartorio (2020: 352) seem to agree). This observation has three significant consequences.

First, this is not a case where we have a clear moral intuition to guide our thinking about graded causation. In this respect, the present cases differ from the cases discussed in the previous section. While it is clear enough that the intuition that dilution of responsibility is to be avoided can guide our thinking about (the relevant notion of) graded causation, it is much less clear what to say about the moral comparison between Victim and Hardy Victim. For this reason, both Bernstein and Sartorio look to Proportionality (and an independently plausible answer to the causal question) for guidance concerning the moral question.

Second, the moral question remains puzzling even if Proportionality is abandoned. That is, abandoning Proportionality does not solve Bernstein's Moral Difference Puzzle, and Proportionality does not create the Moral Difference Puzzle. The consequence of abandoning Proportionality would merely be that we have to look elsewhere for guidance concerning the moral question.

Third, one could see this as a 'spoils to the victor'-situation: since common sense does not give us a clear answer to the moral question, we can, to use David Lewis' words, 'reasonably accept as true whatever answer comes from the analysis that does best on the clearer cases' (Lewis, 1986: 194). In our case, this means our treatment of the cases of redundant causation and joint causation in the previous section can inform our treatment of the present cases. If an answer to the moral question naturally springs out of a theory that can adequately handle dilution of responsibility and that is also consistent with the more general role that graded causation plays in prediction and explanation, then this gives us at least some reason to accept that answer as correct.

That said, let us now turn to the causal question. Bernstein and Sartorio both claim that the causal question is just as puzzling as the moral question. If this is the case, Proportionality cannot guide our treatment of the moral question. However, the argument given in the previous section clearly implies that Causal difference: greater is the correct answer. Each assassin in Victim causally contributes more to the victim's death than each assassin in Hardy Victim. In Victim, the action of one of the assassins is sufficient on its own to kill the victim (given the circumstances); in Hardy Victim, the action of one of the assassins alone is not sufficient to kill the victim. Therefore, an account of graded causation relying on the Sufficiency Criterion implies that the assassins in Victim each make a higher causal contribution than the assassins in Hardy Victim.

Why do Bernstein and Sartorio think that the causal question does not have a determinate answer? Obviously, Victim is a case of symmetric overdetermination, whereas Hardy Victim is a case of joint causation. Bernstein, therefore, correctly assumes that the respective causal contribution of the assassins in Victim is different from the respective causal contribution of the assassins in Hardy Victim. However, she considers it intuitively unclear whether the situation should be evaluated according to the Necessity Criterion or according to the Sufficiency Criterion.

The Necessity Criterion suggests that each assassin's contribution in Hardy Victim is higher than in Victim. On the other hand, the Sufficiency Criterion suggests the opposite, that is, that each assassin's contribution in Victim is higher than in Hardy Victim. Further, according to the argument given in the previous section, this is the criterion on which a theory of graded causation should be based.

Sartorio emphasizes that both criteria influence our pre-theoretic thinking about graded causation. At one point, she offers the following perspective on the matter:

I find myself wanting to say: well, in a sense, your contribution is more significant in the first case, and in another sense, it is more significant in the second case. But this is unhelpful. For what we wanted to know, ultimately, is which of the two agents is more morally responsible for the victim's death in light of the causal contribution they made. In other words, what we need is an account of causal contribution that can be relevant to attributions of moral responsibility. We want to know which agent makes a more significant contribution, in the sense that matters to moral responsibility. And what is that sense? I feel like I do not know how to answer that question. (Sartorio, 2020: 352)

It is helpful to think of the Necessity Criterion and the Sufficiency Criterion as corresponding to two different concepts of graded causation. Call the concept corresponding to the Necessity Criterion 'contribution_{NEC}' and the concept corresponding to the Sufficiency Criterion 'contribution_{SUFF}'. We agree with Sartorio that Proportionality cannot shed light on moral responsibility as long as it is left open whether it relies on contribution_{NEC} or contribution_{SUFF}. That is, we need to clarify the status of the following two versions of Proportionality:

Proportionality_{NEC}: The degree of responsibility of an agent A for an outcome o is (ceteris paribus) proportionate to the degree to which A's actions (or omissions) contributed_{NEC} to o.

Proportionality_{SUFF}: The degree of responsibility of an agent A for an outcome o is (ceteris paribus) proportionate to the degree to which A's actions (or omissions) contributed_{SUFF} to o.

As we have already seen in Sect. 2, Proportionality_{NEC} is highly problematic since it leads to dilution of responsibility and does not allow for any causal (or moral) differentiation between necessary causes. In other words, Proportionality_{NEC} does not do very well on the clearer cases. This should raise suspicion about its applicability to

the comparison between Victim and Hardy Victim. A principle that clearly misleads us in the cases discussed in Sect. 2 is likely to mislead us in the present cases as well.

Proportionality_{SUFF}, on the other hand, does not lead to any obviously false judgments in other cases. On the contrary, it can explain why there is no dilution of responsibility and why not all necessary causes are morally on par. So, before even considering the comparison between Victim and Hardy Victim, it is quite clear that Proportionality_{NEC} is problematic, whereas Proportionality_{SUFF} has a good chance of being a fruitful principle. Contribution_{NEC} is not connected to moral responsibility in the right way and does not adequately account for the function that degrees of causation play in prediction and explanation. This is a good reason to dismiss Proportionality_{NEC} and consider Bernstein's puzzle cases in light of Proportionality_{SUFF} instead.

At first sight, Proportionality_{SUFF} brings with it a clear answer to the Moral Difference Puzzle. In light of Proportionality_{SUFF}, Moral difference: greater is the correct answer. The assassins in Victim are 'more morally responsible in light of the causal contribution they made' (Sartorio, 2020: 352). One could stop here and accept the answer suggested by Proportionality_{SUFF}. This would be in line with the above-mentioned idea that the comparison between Victim and Hardy Victim could be seen as a 'spoils to the victor'-situation. Since the pre-theoretic evaluation of the cases is unclear and Proportionality_{SUFF} does best on the clearer cases, we should accept as true whatever answer comes from Proportionality_{SUFF}.

However, this option remains unsatisfactory in some respects. First, there is still a mismatch between the pre-theoretic evaluation and the evaluation that results from Proportionality_{SUFF}. While Proportionality_{SUFF} strongly favors Moral difference: greater, the pre-theoretic evaluation is unclear. Even Lewis states that 'it would still be better, however, if theory itself went undecisive about the hard cases' (Lewis, 1986: 194). Second, and relatedly, Proportionality_{SUFF} cannot explain the intuitive pull of the contrary resolution of the Moral Difference Puzzle. Bernstein (2017: 172) emphasizes that there is something to be said in favor of 'Moral difference: lesser' as well: 'in a way [...] each assassin is more essential to the victim's death in (Hardy Victim) than in (Victim), for the victim's death would have occurred no matter what the second assassin did in (Victim), but not if the second assassin had opted out in (Hardy Victim)' (Bernstein, 2017: 171). It is certainly true that there is such a difference between the two cases, and it also seems that this difference is morally significant. The crucial point is, nevertheless, that this moral difference is not best explained via a graded notion of causal contribution along the line of the Necessity Criterion or a principle like Proportionality_{NEC}.

To see this, compare Victim to Victim-3:

Victim-3: Three independently employed assassins, unaware of each other, are dispatched to eliminate Victim. Being struck by one bullet is sufficient to kill Victim. Each assassin shoots, and Victim dies.

In Victim-3, Victim's death is even more overdetermined than in Victim. Each assassin's action is further away from being necessary than in Victim. (In Victim, one change to the actual situation would render any assassin's shot necessary, whereas, in Victim-3, two changes are needed to render any assassin's shot necessary.) Still,

intuitively, the assassins are not less responsible than in *Victim*. This, again, is the same insight that dilution of responsibility should be avoided, as previously.

This shows that it is not generally the degree of contribution_{NEC} that is relevant for moral responsibility. Bernstein's observation inherits its plausibility not from a general principle like Proportionality_{NEC}; instead, it seems to be specifically the difference between strictly necessary causes and non-necessary causes that is relevant here. Once a cause is non-necessary, how close it is to being necessary is irrelevant. In other words, the relevance of necessity for moral responsibility is not a matter of degree but rather an all-or-nothing affair.

Why could the question of whether a cause is necessary or not be relevant to moral responsibility? The key point here seems to be that the assassins in *Victim* lack an ability that the agents in *Hardy Victim* have. The assassins in *Hardy Victim* can single-handedly prevent the outcome. It is in this sense that their actions are more essential to the outcome. Plausibly, the ability to prevent an outcome is morally significant. Since the assassins in *Hardy Victim* could have prevented the outcome by omitting their actions, they deserve additional blame for failing to do so.⁹

Agents can have the ability to prevent an outcome even if their actual actions are not necessary for the outcome. In such a case, the ability to prevent an outcome still makes a moral difference. For illustration, consider a variant of *Victim* in which one of the assassins was not only able to refrain from shooting, but also able to prevent the other assassin from shooting.¹⁰ In this case, just like in the original *Victim*-case, the assassin's actual action was sufficient and non-necessary for the victim's death. However, the assassin was also able to prevent the outcome. It seems correct to say that the assassin in this case is even more responsible for the death than in *Victim*. After all, they could have prevented the outcome but chose not to do so. Accordingly, they deserve additional blame for failing to prevent the victim's death. This difference cannot be explained by Proportionality_{NEC}: the assassins' actual actions in the present case and in the original *Victim*-case are equally close to being necessary for the outcome. However, the difference can be explained by the relevance of the ability to prevent the outcome: the assassin has the ability to prevent the outcome in the present case, while they lack this ability in the original *Victim*-case. This further supports the view that the ability to prevent an outcome is morally significant, and that the pull of the idea that the assassins in *Hardy Victim* are more responsible than in *Victim* is best explained by this observation.

However, if there is a morally significant difference between *Victim* and *Hardy Victim* due to the ability to prevent the outcome, then the comparison between these two cases violates the *ceteris paribus* condition in Proportionality_{SUFF} (and also the *ceteris paribus* condition in Proportionality). The assassins in *Victim* and *Hardy Victim* indeed differ with respect to how close their actions are to being sufficient for the outcome. But the two agents also differ concerning their ability to prevent the outcome. Since the latter difference is morally significant, the two cases are not comparable. Comparing them is like comparing two cases in which two agents caused

⁹ The relevance of the ability to prevent an outcome for moral responsibility is also discussed by Sartorio (2005).

¹⁰ Thanks to an anonymous reviewer for suggesting this case.

the same harmful outcome by their actions, but one of them acted freely, whereas the other one was forced to act. The moral evaluation of these cases will be different, independent of whether the causal contribution of the actions was the same.¹¹

This suggests the following analysis of the Moral Difference Puzzle: a theory of graded causation based on the Sufficiency Criterion provides a clear answer to the causal question. Each assassin in Victim causally contributes more to the victim's death than each assassin in Hardy Victim (Causal difference: greater). If the *ceteris paribus* condition were fulfilled, Proportionality_{SUFF} would imply that each assassin in Victim is also more responsible for the victim's death than each assassin in Hardy Victim (Moral difference: greater). However, there is reason to think that the *ceteris paribus* condition is not fulfilled since whether an agent had the ability to prevent an outcome is a morally relevant factor that must be taken into account when assessing their moral responsibility. This explains our intuitive resistance to the claim that the assassins in Victim are more morally responsible for the victim's death than those in Hardy Victim.¹²

An analogous consideration applies to the five pollution cases discussed in Sect. 2. Pollution-1, Pollution-2, and Pollution-3, the three non-redundant scenarios, satisfy the *ceteris paribus* condition in Proportionality_{SUFF}: all agents involved in these cases had the ability to prevent the outcome. However, according to a theory of graded causation relying on the Sufficiency Criterion, their causal contributions differ and therefore, by Proportionality_{SUFF}, their moral responsibility differs as well: in Pollution-2, company A is more responsible for the damage than in Pollution-1 and Pollution-3. If we compare Pollution-4 and Pollution-5, the two redundant scenarios, it should be clear that they also satisfy the *ceteris paribus* condition since none of the agents had the ability to prevent the outcome (and all other relevant factors are held equal by stipulation). Therefore, a theory of graded causation relying on the Sufficiency Criterion implies that all agents make the same causal contribution in both scenarios, which in turn means that they are all morally on par. All agents in both scenarios are equally morally responsible for the death of the fish species.

By contrast, comparisons between redundant and non-redundant scenarios would violate the *ceteris paribus* condition in Proportionality_{SUFF}. Compare, for instance, Pollution-2 to Pollution-4. A theory of graded causation based on the Sufficiency Criterion implies that all agents in both scenarios make the same causal contribution. However, that does not also imply that they are morally on par since they differ with respect to their ability to prevent the outcome.

Finally, this leads to the following diagnosis for the Moral Difference Puzzle: the Moral Difference Puzzle is a genuine puzzle since moral intuitions about Victim and Hardy Victim pull in opposing directions. If one's intuitions are primarily guided by

¹¹ One might object that this kind of reasoning can also be applied to the comparison between Pollution-2 and Pollution-3. Even though in both cases, company A had the ability to prevent the outcome, only in Pollution-2 did it have the ability to produce the outcome on their own. Since this ability is also morally relevant, the *ceteris paribus* condition in Proportionality is violated; therefore, Proportionality does not allow the conclusion that company A in Pollution-2 is more responsible than in Pollution-3. However, there is no reason to assume that, in Pollution-3, company A lacked the ability to produce the outcome on its own. That it dumped only five units does not imply that it did not have the ability to dump ten units.

¹² For discussion of a further intuitive puzzle related to Proportionality, see Demirtas (2022a).

proportionality considerations and the assessment that the causal contributions of the assassins in Victim are greater than in Hardy Victim, one is inclined to conclude that Moral difference: greater is the correct verdict, that is, each assassin in Victim is more morally responsible for the victim's death than each assassin in Hardy Victim. On the other hand, if one's intuitions are primarily guided by the observation that the morally relevant abilities (that is, the ability to prevent the death of the victim) of the agents differ, one is more inclined to conclude that Moral difference: lesser is the correct verdict, that is, each assassin in Victim is less morally responsible for the victim's death than each assassin in Hardy Victim.

However, the Moral Difference Puzzle is not a *causal* difference puzzle. If a theory of graded causation is based on the Sufficiency Criterion, as we have argued, the verdict on the causal contributions of the assassins in the two cases is clear. We should opt for Causal difference: greater since each assassin in Victim makes a greater causal contribution to the victim's death than either assassin in Hardy Victim.

Proportionality_{SUFF} adequately captures these considerations. By relying on a theory of graded causation based on the Sufficiency Condition, it supports a clear causal verdict on the Victim and Hardy Victim cases. By including a *ceteris paribus* condition, it leaves the possibility that the moral verdict could be less clear.

4 Conclusion and a note on moral luck

Let's take stock: we have argued that proportionality considerations lead to two adequacy conditions for accounts of graded causation in the context of discussions about moral responsibility. First, such accounts should avoid dilution of responsibility, and adding independent redundant causes to a causal structure should not diminish the causal contribution of each cause. Second, not all necessary causes are morally on par, and an appropriate account of graded causation should allow for differences in the causal contributions of strictly necessary causes. Accounts of graded causation based on the Necessity Criterion always violate both conditions, whereas accounts of graded causation based on the Sufficiency Criterion can satisfy them. We conclude that a theory of graded causation that captures the relationship between degrees of causation and degrees of moral responsibility should be based on the Sufficiency Criterion.

Starting from this observation, we have defended a version of Proportionality, namely Proportionality_{SUFF}, against the charge of creating confusion about moral responsibility. We have shown how Proportionality_{SUFF} allows for a plausible analysis of the puzzling comparison between Victim and Hardy Victim, discussed by both Bernstein (2017) and Sartorio (2020). According to this analysis, the assassins in the two cases differ with respect to their causal contribution. Moreover, they also differ with respect to their abilities to prevent the outcome, which is another morally relevant factor. This explains why we find it intuitively difficult to compare the two cases in terms of moral responsibility.

What has yet to be considered is the role of this argument in the context of the debate on moral luck. Bernstein argues that Proportionality gives rise to an interesting new form of resultant moral luck (Bernstein, 2017: 167–170). Standard cases of

moral luck occur when factors entirely outside an agent's control make a difference in the moral evaluation of the agent. For example, two otherwise identical drivers deserve a different moral evaluation if one kills a pedestrian while the other does not, even if the presence or absence of the pedestrian is entirely out of the drivers' control (Nagel, 1976). Proportionality luck arises whenever factors outside an agent's control make a difference in the agent's *degree* of moral responsibility, where the difference in the degree of moral responsibility is explained by a difference in the degree of causal contribution.

That there are cases of proportionality luck is compatible with the argument of this paper. Reconsider the scenario briefly discussed in Sect. 2, where company B's action has a causal influence on company A's action. In this scenario, company A dumps ten units of chemicals, and under normal circumstances, this would be sufficient to kill the fish species. However, company A is lucky since company B also discharges chemicals, and the chemicals released by company B partially neutralize the effect of the chemicals dumped by company A. Therefore, if company B discharges these chemicals, company A's action is less close to being sufficient than in the case where company B does nothing, and Proportionality_{SUFF} implies that company A's responsibility is reduced – even though company B's action was entirely outside their control.

Obviously, a detailed discussion of whether this commitment to proportionality luck is problematic is beyond the scope of this paper. Suffice it to point out that noting that more needs to be said about moral luck and the cases leading to it only suggests one possible direction in which the debate could go – not that the theory of graded causation and proportionality proposed in this paper is flawed.

Acknowledgements This paper was presented at the online summer school “Causation and Responsibility” at the University of Bern in 2021 and at the colloquium on Philosophy of Mind and Metaphysics at the University of Bern in 2022. We would like to thank the audiences at these events for helpful comments and discussion. Special thanks go to Huzeyfe Demirtas, Nicole Dolby-Rathgeb, Kian Salimkhani, Jonas Werner, and two anonymous referees for extremely valuable input and feedback on various parts of the argument, and to Nino Töndury for careful and competent linguistic corrections of the manuscript.

Funding Open access funding provided by University of Bern

Statements and Declarations

The research leading to these results received funding from the Swiss National Science Foundation under Grant Agreement No. 100019E_189589.

The authors have no relevant financial or non-financial interests to disclose.

Both authors contributed equally to the research and to the preparation of the manuscript.

Both authors read and approved the final manuscript.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Bernstein, S. (2017). Causal proportions and moral responsibility. In D. Shoemaker (Ed.), *Oxford Studies in Agency and Responsibility* (Vol. 4, pp. 165–182). Oxford University Press.
- Braham, M., & Hees, M. (2009). Degrees of causation. *Erkenntnis*, 71(3), 323–344.
- Brun, G. (2016). Explication as a method of conceptual re-engineering. *Erkenntnis*, 81(6), 1211–1241.
- Chockler, H., & Halpern, J. Y. (2004). Responsibility and blame: A structural-model approach. *Journal of Artificial Intelligence Research*, 22, 93–115.
- Demirtas, H. (2022a). Causation comes in degrees. *Synthese*, 200(1), 1–17.
- Demirtas, H. (2022b). Moral responsibility is not proportionate to causal responsibility. *Southern Journal of Philosophy*, 60(4), 570–591.
- Dowe, P. (2000). *Physical Causation*. Cambridge University Press.
- Ferey, S., & Dehez, P. (2016). Overdetermined causation cases, contribution and the Shapley Value. *Chicago-Kent Law Review*, 91(2), 637–658.
- Gill, M., Kominsky, J., Icard, T., & Knobe, J. (2022). An interaction effect of norm violations on causal judgement. *Cognition*, 228, 105–183.
- Godfrey-Smith, P. (2009). Causal pluralism. In H. Beebe, P. Menzies, & C. Hitchcock (Eds.), *The Oxford Handbook of Causation* (pp. 326–337). Oxford University Press.
- Hall, N. (2004). Two concepts of causation. In J. Collins, N. Hall, & L. A. Paul (Eds.), *Causation and Counterfactuals* (pp. 181–276). MIT Press.
- Halpern, J., & Hitchcock, C. (2015). Graded causation and defaults. *British Journal for the Philosophy of Science*, 66(2), 413–457.
- Hart, H. L. A., & Honoré, T. (1959). *Causation in the Law*. Oxford University Press.
- Kaiserman, A. (2016). Causal contribution. *Proceedings of the Aristotelian Society*, 116(3), 387–394.
- Kaiserman, A. (2017). Partial liability. *Legal Theory*, 23(1), 1–26.
- Kaiserman, A. (2018). ‘More of a Cause’: Recent work on degrees of causation and responsibility. *Philosophy Compass*, 13(7), e12498.
- Kaiserman, A. (2021). Responsibility and the ‘Pie Fallacy’. *Philosophical Studies*, 178(11), 3597–1616.
- Lewis, D. (1973). Causation. *Journal of Philosophy*, 70, 556–567.
- Lewis, D. (1986). Postscripts to ‘Causation’. In D. Lewis (Ed.), *Philosophical Papers* (Vol. II, pp. 172–213). Oxford University Press.
- Moore, M. S. (2009). *Causation and Responsibility: An Essay in Law, Morals, and Metaphysics*. Oxford University Press.
- Nagel, T. (1976). Moral luck. *Aristotelian Supplementary Volume*, 50(1), 115–152.
- Nelkin, D. (2016). Difficulty and degrees of moral praiseworthiness and blameworthiness. *Noûs*, 50(2), 356–378.
- Northcott, R. (2013). Degree of explanation. *Synthese*, 190(15), 3087–3105.
- Psillos (2007). Causal pluralism. In C. Gershenson, D. Aerts, & B. Edmonds (Eds.), *Worldviews, Science and Us: Philosophy and Complexity* (pp. 131–151). World Scientific Publishing Co Pte Ltd.
- Quillien, T., & Barlev, M. (2022). Causal judgement in the wild: Evidence from the 2020 U.S. Presidential Election. *Cognitive Science*, 46, e13101.
- Sartorio, C. (2005). A new asymmetry between actions and omissions. *Noûs*, 39(3), 460–482.
- Sartorio, C. (2020). More of a cause? *Journal of Applied Philosophy*, 37(3), 346–363.
- Schaffer, J. (2003). Overdetermining causes. *Philosophical Studies*, 114(1–2), 23–45.
- Sprengr, J. (2018). Foundations of a probabilistic theory of causal strength. *Philosophical Review*, 127(3), 371–398.
- Tadros, V. (2018). Past killings and proportionality in war. *Philosophy and Public Affairs*, 46(1), 9–35.
- Zimmerman, M. (1985). Sharing responsibility. *American Philosophical Quarterly*, 22(2), 115–122.