



Three Kinds of Causal Indeterminacy

Vera Hoffmann-Kolss

To cite this article: Vera Hoffmann-Kolss (04 Aug 2024): Three Kinds of Causal Indeterminacy, Australasian Journal of Philosophy, DOI: [10.1080/00048402.2024.2381120](https://doi.org/10.1080/00048402.2024.2381120)

To link to this article: <https://doi.org/10.1080/00048402.2024.2381120>



© 2024 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group



Published online: 04 Aug 2024.



Submit your article to this journal [↗](#)



Article views: 119



View related articles [↗](#)



View Crossmark data [↗](#)



Three Kinds of Causal Indeterminacy

Vera Hoffmann-Kolss 

University of Bern

ABSTRACT

The goal of this paper is to argue that there is indeterminacy in causation. I present three types of cases in which it is indeterminate whether an event c caused another event e : (1) cases of absence causation recently discussed by Bernstein and by Swanson, (2) cases leading to Sorites paradoxes for causation, and (3) cases where c and e occur in certain indeterministic causal structures and it is therefore indeterminate whether there is a causal relation between them. These cases, I argue, provide strong evidence that indeterminacy is an important general feature of the causal relation and that philosophical theorizing about causation should take this observation into account.

ARTICLE HISTORY Received 23 July 2022; Revised 7 July 2023

KEYWORDS causation; probabilistic causation; absence causation; prevention; Sorites paradox; indeterminacy

1. Introduction

One hour before the start of the 10am classes, the janitor of a large university building learns that the lighting systems in four lecture rooms do not work. Fixing one of the systems takes almost an hour, so the most he can do is fix one of them. Not feeling up to discussing with the lecturers which of the systems he should fix, he decides to stay in his office and do nothing. Consequently, four 10am classes must be cancelled.

Cases such as this one have recently attracted attention in the philosophical debate on causation (Sartorio 2006; Ballarín 2014; Bernstein 2016a; Swanson 2017). Suppose that Jane was going to teach one of the classes that had to be cancelled. Was the janitor's doing nothing a cause of Jane having to cancel her class? Trying to answer this question leads to a dilemma. Since the janitor could have fixed at most one of the systems, it is implausible that his doing nothing caused all four classes to be cancelled. However, if his doing nothing were considered a cause of Jane's having to cancel her class, then parity of reasoning would imply precisely that: his doing nothing was also a cause that the other three lecturers had to cancel their classes, that is, that all four classes had to be cancelled.¹ Therefore, assuming that the janitor's doing nothing caused Jane to cancel her class leads to a contradiction.

¹This holds at least if causation is additive in the following sense: if c causes e_1 , and c causes e_2 , and \dots , and c causes e_n , then c causes $(e_1 \& e_2 \& \dots \& e_n)$, that is, the conjunction of e_1, e_2, \dots, e_n . Most philosophers of

On the other hand, assuming that the janitor's doing nothing was not a cause of Jane's having to cancel her class leads to an equally implausible result. By parity of reasoning, his doing nothing was not a cause of any of the other classes having to be cancelled. But given that he would have been able to fix one of the systems, this also seems implausible.

Bernstein and Swanson both argue that this puzzle can be solved in the following way: even though the janitor's doing nothing caused *one* lecturer to cancel her class, it is indeterminate which one (Bernstein 2016a; Swanson 2017; see Ballarin 2014 for a similar diagnosis). Therefore, the question of whether the janitor's doing nothing caused Jane to cancel her class does not have a determinate answer, and causal claims are sometimes indeterminate.²

If causal statements can indeed be indeterminate, this will have consequences for the way in which we understand and use the concept of causation. From a metaphysician's point of view, the claim that it may be indeterminate whether *c* caused *e* raises the question of whether causal indeterminacy is a particular type of worldly indeterminacy. If one thinks that causation is a relation in the world that exists independently of being described or conceptualized, and if there are scenarios in which it is in principle impossible to decide whether *c* caused *e*, then this seems to show that the world is indeterminate in a certain respect.

Moreover, causal indeterminacy can lead to indeterminacy in other areas. For instance, it is plausible that there is a close connection between causation and moral responsibility: an agent can only be held morally responsible for a (possibly harmful) state of affairs, such as the occurrence of an accident, if she causally contributed to it (Bernstein 2017; Kaiserman 2021; for discussion see, however, Sartorio 2004). But if causal contribution is a necessary condition for moral responsibility and if it is indeterminate whether an agent's action caused a certain outcome, then it is indeterminate whether the agent is morally responsible for the outcome—even if all other preconditions for moral responsibility (such as not having been forced to perform the action, having been in a position to know about the consequences of the action, etc.) are fulfilled. This may have consequences for our understanding of the concept of moral responsibility.

So far however, the question of whether it can be indeterminate whether *c* caused *e* has received comparatively little attention. The only cases that have been discussed in detail are cases having the structure of the janitor example described above. As systematically important as they are, these cases are also problematic in that they rely on the controversial assumption that absences or omissions can be causes. If this assumption is rejected, the indeterminacy disappears: the janitor's doing nothing does not qualify as a cause of Jane having to cancel her class, since all causal statements involving absences are determinately false. (The view that absences cannot be causes is held, for example, by Dowe 2000; Beebe 2004.) Moreover, even if one agrees with authors who hold that absence causation is possible (Woodward 2003; Schaffer 2004; Bernstein 2014, 2016b), the claim that causal claims can be indeterminate will

causation seem to accept this principle. Sartorio (2006) argues, however, that cases like the janitor example are a reason to reject it.

²This does not imply that all causal statements that can be made about this situation are indeterminate. For instance, the statement 'the janitor's doing nothing caused someone to cancel their class' seems to be determinately true. This is not problematic, since the example is only meant to show that there are indeterminate causal statements, not that causal indeterminacy holds across the board.

be considerably stronger, if one can show that indeterminacy occurs in very different causal structures. The goal of this paper is to show just that.

More specifically, I argue that, in addition to the type of causal indeterminacy arising from cases having the structure of the janitor example, there are two other types of causal indeterminacy. A second type of causal indeterminacy arises because some causal statements allow the construction of Sorites paradoxes. If it is possible to construct a Sorites series for causal relations, then there will be borderline cases of causation, and these borderline cases are instances of causal indeterminacy (for a more detailed discussion of the relationship between Sorites susceptibility, the possibility of borderline cases, and indeterminacy, see Keefe 2000; Barnes 2010). I will describe two kinds of causal structure leading to a Sorites paradox. One of them also relies on the assumption that there can be causation by omission. The other one is an overdetermination scenario, which does not involve causation by omission.

The third possibility to show that causal statements can be indeterminate relies on the observation that there are structures in which indeterministic processes lead to causal indeterminacy: I describe three scenarios—a prevention, an overdetermination, and a trumping scenario—in which this is true. These cases are quite different from the previous cases, since they suggest that causal claims are indeterminate, but do not lead to a Sorites paradox.

The next section is devoted to developing two Sorites paradoxes for causation (Section 2). I then argue that causal statements can be indeterminate if they describe relations that occur in structures containing indeterministic relations or processes (Section 3). I conclude by discussing two possible objections to my argument (Section 4).

Before moving on to the next section, it should be noted that the cases of causal indeterminacy discussed here are not the only indeterminacy phenomena related to causation. Torrago defines ‘vague causation’ as ‘a causal relation that holds amongst events at least one of which is vague’ (Torrago 2000: 315). Accordingly, the claim that Jimmy’s baldness causes his dissatisfaction with his appearance may be indeterminate if it is indeterminate whether Jimmy is bald. A discussion of this phenomenon would be interesting in itself but is beyond the scope of this paper. This paper is concerned with causal statements that are indeterminate *even though* they hold between precise events, that is, between events for which it is determinate whether they occurred.

Sometimes, the term ‘causal indeterminacy’ is also used in cases in which an effect does not have a sufficient cause (Ellis 2004: 331). This use is different from the way in which the term is used here. The question discussed in this paper is not whether there are causes that fail to determine their effects. It may be determinately true that Jimmy’s lung cancer was caused by his heavy smoking. However, that does not mean that there was a deterministic relationship between his smoking habits and his disease. Conversely, the cases of causal indeterminacy having the structure of the janitor example discussed above do not involve indeterministic relationships at all. The thesis of this paper is that there are cases in which it is indeterminate whether there is a causal relation between two events, not that indeterministic relations can be causal.³

³ The claim that there is indeterminacy in causation should also be distinguished from the claim that causation allows for degrees, that is, that if events c_1 and c_2 both cause a further event e , it is possible that c_1 is more

2. Two Sorites Paradoxes for Causation

Some properties and relations are vague. *Being bald* and *being tall* are classical examples of vague properties, *being much taller than* and *being in the vicinity of* are examples of vague relations. As pointed out above, a central characteristic of vague notions is that they lead to Sorites paradoxes. For instance, if d measures the difference in height between Jimmy and Jane, a Sorites paradox for the *much-taller-than* relation can be constructed as follows:

- (1) If $d = 0.5$ inches, then Jimmy is not much taller than Jane.
- (2) For any number x , if Jimmy is not much taller than Jane when $d = x$ inches, then Jimmy is not much taller than Jane when $d = x + 0.01$ inches.

Therefore:

- (3) If $d = 20$ inches, then Jimmy is not much taller than Jane.

The paradox consists in the fact that (3) follows validly from (1) and (2), but (1) and (2) are intuitively true, whereas (3) is intuitively false.

To see that intuitive judgments about causal structures can also lead to Sorites paradoxes, consider a further causation-by-omission scenario:

Omissions

Jimmy and his family have several plants and plan to go on holiday. A couple of days before they leave, Jimmy's mother-in-law comes by and gives them an irrigation system that they can use to water the plants during their absence.

Case 1

The irrigation system is very reliable. It has already worked on multiple occasions, and Jimmy has no problems installing it correctly. However, contrary to what was to be expected, the system fails from the first day of the holiday, and the plants die.

Case 2

The irrigation system is very unreliable. It has already failed on multiple occasions, and Jimmy installs it only to avoid the argument with his mother-in-law. As expected, the system fails, and the plants die.

Did the irrigation system's failure cause the plants to die? In case 1, the answer seems to be 'yes'. The system was installed correctly and could be expected to function properly. No additional measures were necessary. Case 2 is different. The system could not be expected to function properly. In this case, it seems much more natural to say that Jimmy's failure to find an alternative solution, for instance, his failure to ask his neighbour to water the plants, caused the plants to die. But then, the scenario leads to a first Sorites paradox for causation (where P_F is the probability of the system's failing):

of a cause of e than c_2 is. Degrees of causation receive increasing attention in the philosophical debate on causation (Demirtas 2022; Kaiserman 2018). But the claim that causation allows for degrees is logically independent of the claim that causal statements can be indeterminate. For instance, the janitor example is meant to show that there is indeterminacy in causation without presupposing degrees of causation. Conversely, one can consistently hold that c_1 and c_2 make different causal contributions to e , even though c_1 is determinately a cause of e and c_2 is determinately a cause of e .

- (1) If $P_F = 0.001$, then the failure of the system caused the plants to die.
- (2) For any number x , if the failure of the system caused the plants to die when $P_F = x$, then the failure of the system caused the plants to die when $P_F = x + 0.001$.

Therefore:

- (3) If $P_F = 0.999$, then the failure of the system caused the plants to die.

Premise (1) describes case 1, in which it was highly likely that the system would function properly, and its unexpected failure caused the plants to die. Premise (2) relies on the intuition that if we have already judged that the failure of the system caused the plants to die, then this judgment will not be altered if we diminish the probability of the system's functioning just a tiny bit. The conclusion (3), which validly follows from (1) and (2), describes case 2, in which the probability of the failure of the system is very high, and contradicts the intuition that if the system could not be expected to function properly, its failure did not cause the plants to die. To avoid this contradiction, there would have to be a transition in the Sorites series from the truth of the causal claim to the falsity of the causal claim. However, since there is no exact cut-off value for P_F that marks this boundary, there are borderline cases where it is indeterminate whether the failure of the system caused the plants to die.

The scenario is, of course, a variation on standard examples of causation by omission. Compare the case in which Jimmy's neighbour, who has previously committed herself to watering the plants, fails to do so with the case in which Hillary Clinton, who is in no way acquainted with Jimmy or his family, fails to water the plants. Intuitively, only the neighbour's failure to water the plants, but not Hillary Clinton's failure to do so should be considered a cause of the plants' death. This judgment is usually justified by moral considerations: if Jimmy's neighbour promises to water the plants, she will have a moral obligation to do so. Hillary Clinton, by contrast, has no moral obligation to water the plants. I will leave open here whether a case similar to the irrigation system example could be constructed for the standard flower watering scenario, since it is controversial whether moral considerations should influence causal judgments in this way.

As pointed out above, an example involving causation by omission can be considered problematic, since it is controversial whether absences can be causes. Moreover, the intuition that whether the failure of the system caused the plants' death depends on how likely it was to fail may not be shared by everyone. However, not all cases leading to a Sorites paradox for causation rely on these controversial assumptions. A further Sorites series arises from the distinction between causal overdetermination and preemption:

Overdetermination vs. pre-emption

Jimmy has a terrible headache and needs a painkiller. Taking one headache pill would be sufficient to cure his pain. Jimmy, not knowing what the correct dosage is, takes two pills, and his headache disappears.

Case 1

Jimmy swallows one of the pills first and the other one five seconds later.

Case 2

Jimmy swallows one of the pills first and the other one four hours later (because he mistakenly believes that if he swallows just one pill, his headache will return).

It should be uncontroversial that in both cases, Jimmy's taking the first pill was a cause of his getting rid of his headache. But was the second pill also a cause of the headache's disappearing? In Case 2, the answer seems to be 'no'. The case has the structure of a late preemption scenario. After four hours, the first pill will have been effective. Accordingly, potential effects of the second pill will be causally preempted by the effects of the first pill.

In case 1, however, both pills relieve Jimmy's pain. The time that elapses between Jimmy taking the first pill and the second pill is so short that it is impossible to determine which of the two was causally relevant in curing his pain. Accordingly, the case has the structure of an overdetermination scenario: there are two causes, each of which would have been sufficient to produce the effect, and neither is privileged over the other.

An immediate objection is that the pill that was taken first is privileged over the second one. Consequently, Jimmy's taking the two pills within a five-second time interval does not count as a case of perfectly symmetrical overdetermination. Only cases in which two (or more) sufficient causes occur at *exactly* the same time should be considered cases of perfectly symmetrical overdetermination.

This would imply, however, that examples typically regarded as standard cases of overdetermination would not work either. Consider, for instance, the case of a victim's being hit by several deadly bullets at the same time. Suppose that the idea is that the victim is hit simultaneously by two bullets if the two bullets reach the surface of her body at exactly the same time. Since the physical boundaries of human beings are vague, there will be no exact time when a bullet reaches a human's body. But if there is no exact time when the first bullet reached the victim's body, it will also be indeterminate whether the second bullet reached her body at exactly the same time. And, of course, the same problem occurs if we assume that the bullets must reach the victim's vital organs, such as her heart or her brain, at exactly the same time.

An analogous consideration holds for the headache pill case. There is no such thing as swallowing two pills at *exactly* the same time. Even if Jimmy guzzles them down in one gulp, it is plausible to assume that one of the pills will reach his stomach slightly earlier than the other one (and, of course, the boundaries of Jimmy's stomach are fuzzy, too). If two events occur so shortly after each other that it is indeterminate which one of them brought about the effect, both should be considered causes of the effect. The scenario in which Jimmy takes one pill first and the other one five seconds later is such a case.

Therefore, if d measures the time delay between Jimmy's swallowing the first and the second pill, a second Sorites paradox for causation can be constructed as follows:

- (1) If $d = 5$ seconds, then both pills were causes of the disappearance of Jimmy's headache.
- (2) For any number x , if both pills were causes of the disappearance of Jimmy's headache when $d = x$ seconds, then both pills were causes of the disappearance of Jimmy's headache when $d = x + 1$ seconds.

Therefore:

- (3) If $d = 14400$ seconds (four hours), then both pills were causes of the disappearance of Jimmy's headache.

Premise (1) captures the overdetermination intuition invoked by case 1, whereas the conclusion (3) contradicts the preemption intuition underlying case 2. Premise (2) relies on the intuition that slightly increasing the period of time between Jimmy's taking the first and the second pill does not make a difference to whether his taking the second pill should be considered a cause of his recovery. It should be noted that it is not possible to reject premise (2) on the grounds that there is a threshold determining when the second pill ceases to be causally effective. For instance, claiming that the second pill ceases to be a cause of Jimmy's recovery when the first pill has started to become effective does not help, since there is no exact time when the first pill starts being effective. Likewise, claiming that the potential effects of the second pill are preempted by the effects of the first pill once Jimmy's headache has disappeared completely is problematic, since the temporal boundaries of headaches are also vague.

There are three possible objections to this scenario. The first is that it relies on an overly liberal notion of overdetermination. If only cases in which the causes of an effect occur at exactly the same time count as overdetermination, and if putative standard cases of overdetermination do not satisfy this strict criterion, this only shows that the standard cases are not cases of genuine overdetermination. It does not show that the notion of overdetermination should be understood in a more liberal way (for arguments expressing general scepticism about the notion of overdetermination, see Lewis 1973: 567, n. 12; Bunzl 1979).

However, such a strict interpretation of the notion of overdetermination creates another problem. Suppose that an event c_1 by itself would be sufficient to produce some effect e , that another event c_2 by itself would also be sufficient to produce e , and that e has no other sufficient causes. If it is indeterminate exactly when c_1 and c_2 occur, it may also be indeterminate whether c_1 occurs later than, earlier than, or simultaneously with c_2 . But if the temporal order of c_1 and c_2 is indeterminate, then c_1 and c_2 are causally on a par, that is, either they both cause e , or neither of them causes e . Assuming that neither c_1 nor c_2 causes e implies that e occurs without having a sufficient cause (since, by presupposition, e has no other sufficient causes). But the conclusion that e is uncaused is very implausible, since e even has two sufficient causes. The remaining option is to assume that c_1 and c_2 both cause e . But this implies the liberal notion of overdetermination underlying the headache pill example.

The second possible objection is that the scenario does not show that the causal relation itself is vague, but only that causal claims can be vague by virtue of involving vague relata. That is, that the scenario leads to a type of causal indeterminacy that I said at the end of the previous section I would not discuss. However, there is a crucial difference between the type of indeterminacy mentioned at the end of the previous section and the type of indeterminacy that follows from the headache pill scenario. Consider again the claim that Jimmy's baldness causes his dissatisfaction with his appearance. If it is indeterminate whether Jimmy is bald, then it may be indeterminate whether this claim is true, since the truth of a causal claim requires that the relata of the causal relation actually occur. In the headache pill example, on the other hand, it is not indeterminate whether the events in question actually occurred (they definitely did), it is

just indeterminate exactly *when* they occurred. One might still worry whether this means that the indeterminacy of the causal relation is just due to, or inherited from, the indeterminacy of the time of the occurrence of the relata (for a related consideration, see Keefe 2000: 14–15). But this would not be problematic. The temporal order of the relata is often crucial for causal relations. The headache pill example shows that the causal structure of a scenario can depend on its temporal structure. Since temporal distance is a gradual notion, there may be borderline cases that leave the causal structure indeterminate. This kind of causal indeterminacy is clearly different from the kind of indeterminacy mentioned at the end of section 1.

The third possible objection is related to the very presupposition that events can have vague temporal boundaries and possibly do not occur at an exact time. One might wonder whether this presupposition makes the argument circular, since one must presuppose indeterminacy in order to argue for causal indeterminacy. Note, however, that there are distinct kinds of indeterminacy involved here: the argument for causal indeterminacy rests on the assumption that the exact temporal location of events can be indeterminate, but not on the assumption that there is causal indeterminacy (which would of course be circular). Still, the argument would be problematic if the assumption that the temporal location of events can be indeterminate were rejected. I will return to this point in section 4.

3. Indeterministic Relations and Causal Indeterminacy

A further case that might lead to causal indeterminacy is the following (slightly adapted) prevention scenario discussed by Collins (2000):

Prevention

Jimmy and his wife, Jane, play soccer in the garden with their eight-year-old daughter, Suzy. Their neighbour owns a precious amphora residing in the corner of her garden. Suzy kicks the ball in the direction of the amphora. Jimmy throws himself into the trajectory of the ball and catches it.

Case 1

Jane was standing nearby. Had Jimmy not caught the ball, she would have caught it. Had they both failed to catch it, it would have hit the amphora.

Case 2

There is a massive brick wall between Jimmy and Jane's and the neighbour's garden. If Jimmy had not caught the ball, it would have hit the brick wall, not the amphora.

Did Jimmy prevent the amphora from being hit? In case 1, the answer seems to be 'yes'. Either he or Jane had to take action to stop the ball, and Jimmy was the one who did, whereas Jane did nothing. In case 2, the answer seems to be 'no'. If Jimmy had not caught the ball, it would have bounced off the wall without doing any damage.

According to Collins (2000), case 1 can be distinguished from case 2 by considering how far-fetched it is to assume that the ball would have hit the amphora had Jimmy not caught it. In case 1, this possibility does not seem far-fetched. If Jane had been standing two yards further away or been distracted for a moment, she would not have been able to catch the ball. It seems a lot more far-fetched, however, to assume that the brick wall in case 2 would have spontaneously toppled the moment the ball flew toward it or that

Suzy's kick would have been hard enough to go through it. Collins suggests that, in general, an event e is causally dependent on an event c (prevention being a type of causal dependence) if e is either counterfactually dependent on c or would be counterfactually dependent on c if some true proposition were 'false in some not too far-fetched way' (Collins 2000: 229). He points out that being far-fetched is a vague notion, and that therefore, his proposed analysis of causation is also vague (Collins 2000: 229).

Prima facie, Collins' argument can be seen as further evidence that causation is a vague relation. But consider the following scenario:

Case 3

If Jimmy had not caught the ball, the ball would have failed to hit the amphora, but only by a quarter of an inch. (Assume that Jane is not around in this case.)

In case 3, a scenario in which Jimmy's action would have made a difference is not very far-fetched. Had Suzy shot the ball from a slightly different angle, the ball would have been on a direct path to the amphora. This possibility does not seem to be more far-fetched than the possibility that Jane, who is standing nearby, does not manage to catch the ball. In case 3, however, Jimmy did not prevent the amphora from being hit. Even if nobody had intervened, the ball would have missed the vase without doing any damage. Accordingly, the difference between case 1 and case 3 does not depend on how far-fetched it is to assume that the ball would have hit the amphora had Jimmy not caught it, but on the consideration that in case 1, there had to be an intervention to prevent the damage, whereas no such intervention was needed in case 3. It follows that the question of whether Jimmy prevented the damage does not primarily depend on far-fetchedness considerations, but on whether any interventions were required to save the amphora, and the latter question usually has a determinate answer. Therefore, it is questionable whether the prevention scenario discussed by Collins provides further evidence for the vagueness of the causal relation.⁴

However, prevention scenarios do lead to a kind of causal indeterminacy if they involve indeterministic processes or relations. For instance, consider a version of the above prevention scenario where the trajectory of the ball is indeterministic:

Indeterministic prevention

Suzy kicks the ball in the direction of the neighbour's amphora (and Jane is not around). The ball takes course to the amphora. However, its exact trajectory is indeterministic, and the probability of it hitting the amphora if Jimmy does not intervene is 0.97.

Did Jimmy prevent the ball from hitting the amphora? At the time when he caught the ball it was indeterminate whether the ball would hit the amphora, and the probability that it would hit the valuable vase was 0.97. Therefore, the probability that Jimmy prevented the ball from hitting the amphora is also 0.97, and the claim that Jimmy prevented the amphora from being hit is neither determinately true nor determinately false.

⁴ Moreover, intuitions about such cases seem to be unstable. McDermott, who developed the original version of this example, suggests that the person catching the ball prevents it from hitting a valuable object even if the ball would have otherwise hit a solid brick wall (McDermott 1995: 525; see also Woodward 2003: 86–87, for discussion).

The general consideration underlying this judgment is that an event can prevent another from occurring only if the latter would have occurred if the former had not occurred. This is a necessary condition for prevention. Since it is indeterminate whether the ball would have hit the amphora if no one had intervened, it is indeterminate whether Jimmy's action satisfies this necessary condition (for an understanding of probability that supports this conclusion, see Jeffrey 1992: 193). Therefore, it is also indeterminate whether Jimmy prevented the amphora from being hit.

A possible objection is that the notion of causal prevention can be understood differently. One might hold that an event prevents another from occurring only if the latter would determinately have occurred if the former had not occurred. According to this reading, Jimmy does not prevent the amphora from being hit, because it is not the case that the ball would determinately have hit the amphora if he had done nothing.

However, such an understanding of prevention implies that the indeterministic prevention scenario is equivalent to a scenario in which it is impossible for the ball to hit the amphora, for example, because the amphora is 3 miles away. This seems problematic, especially if the probability that the ball would have hit the amphora is further increased so that it is almost certain that the amphora would have been shattered but for some tiny remaining chance. There may be a notion of prevention according to which there is no difference between this case and the case where the amphora is 3 miles away, but it does not seem to be the most natural one.

However, the conclusion that it is indeterminate whether Jimmy prevented the amphora from being shattered could also be undermined if the notion of prevention is understood in the following way: consider the proposition that COVID-19 vaccinations prevent infections. This proposition is determinately true, even though it is usually indeterminate how many infections would have occurred in a group of people if they had not been vaccinated. By analogy, one might say that Jimmy determinately prevented the amphora from being smashed, even though it is indeterminate whether the ball would have hit it. But this analogy ignores the distinction between type and token causation. The indeterminacy thesis defended in this paper applies to token causation, that is, causal relations between concrete events: it is sometimes indeterminate whether an event c caused an event e , even though it is determinately true that both c and e occurred. The indeterminacy of a token causal claim may be compatible with the determinate truth of the corresponding type causal claim. It may be indeterminate whether Jimmy's third COVID-19 vaccination prevented him from getting COVID-19, even though it is determinately true that COVID-19 vaccinations prevent infections. By analogy, it may be indeterminate whether Jimmy's catching Suzy's ball prevented the amphora from being hit (token causation), even though it is determinately true that catching balls that go astray generally prevents accidents (type causation).

It cannot be ruled out that there are further alternative readings of the notion of prevention, according to which it is not indeterminate whether Jimmy prevented the ball from hitting the amphora. The indeterministic prevention scenario shows, however, that there is at least one natural reading that leads to causal indeterminacy.

A more general reason to be sceptical of the indeterministic prevention scenario is that, like the irrigation system example, it presupposes that absences can occur in causal relations, in this case, that the effect can be an absence (the ball not hitting

the amphora). Thus, if the assumption that absences can be the relata of the causal relation, be it on the cause or on the effect side, is rejected, the example becomes problematic. However, there are other causal structures that do not involve absences, in which indeterministic relations also give rise to causal indeterminacy. Consider the following overdetermination scenario:

Indeterministic overdetermination

Jane is a good hacker and hacks her e-mail program so that it makes a sound when it receives a message from an address that starts with the letter 'a'. To keep the program from making too much noise, she also builds in a random process that ensures that the probability of the program making a sound is only 0.1. One day, Jane receives emails from Aurelia and August, whose email addresses both begin with the letter 'a', at exactly the same time. The program does indeed make a sound.

Did Aurelia's e-mail cause the sound? Since the relationship between incoming e-mails and the sound involves a random process, this question does not have a determinate answer. The probability that Aurelia's message caused the sound is 10/19, the probability that August's message caused the sound is likewise 10/19, and the probability that both messages caused the sound is 1/19.⁵ Put another way, a necessary condition for Aurelia's message being a cause of the sound is that the sound was not caused by August's message alone. The probability that the sound was caused by August's message alone is 9/19.⁶ Accordingly, it is indeterminate whether this necessary condition on the causal relevance of Aurelia's message is satisfied, and the probability that it is satisfied is $1 - 9/19 = 10/19$.

A possible objection is that even though the sound is triggered by a random process, so that it is indeterminate before one of the messages is sent whether the sound will occur, it would be possible to examine the internal mechanism of the computer afterwards and to determine whether Aurelia's message, August's message, or both caused the sound. This type of objection is avoided by the next scenario.

This scenario is a variation of the trumping preemption scenarios developed by Schaffer. One of Schaffer's examples is set in a magic world governed by deterministic laws. In this world, the enchantment at midnight is determined by the first spell cast on the day. If both Merlin and Morgana cast spells to turn the prince into a frog on a given day, but Merlin's spell came first, then it was Merlin's but not Morgana's spell that caused the transformation of the prince (Schaffer 2000: 165). Now, consider an indeterministic version of this scenario:

Indeterministic trumping

The laws of magic are indeterministic. If there is only one spell cast on a particular day, then this spell will determine what happens at midnight. However, if there are two (or more) spells cast on a particular day, then the probability that the first spell determines what happens at midnight is 0.9, and the probability that the second, but not the first spell determines what

⁵ Define P_S as the probability of the program's making a sound, P_1 as the probability that the program makes a sound after receiving a message from Aurelia, and P_2 as the probability that the program makes a sound after receiving a message from August. Given that $P_1 = P_2 = 0.1$ (and it is plausible that there is no interaction between them), $P_S = 1 - 0.9 \times 0.9 = 0.19$. Therefore, the probability that Aurelia's message caused the sound is $P_1 / P_S = 10/19$, the probability that August's message caused the sound is $P_2 / P_S = 10/19$, and the probability that both messages caused the sound is $P_1 \times P_2 / P_S = 1/19$.

⁶ $p(\text{sound was caused by August's message alone}) = p(\text{sound was caused by August's message}) - p(\text{sound was caused by both Aurelia's and August's messages}) = P_2 - P_1 \times P_2 / P_S = 10/19 - 1/19 = 9/19$ (see also fn. 5).

happens at midnight is 0.1 (and further spells are causally inert). Both Merlin and Morgana cast spells to turn the prince into a frog. Merlin casts the first spell of the day, Morgana casts the second spell. The prince is turned into a frog at midnight.

Was Merlin's spell a cause of the prince's being turned into a frog? If the laws of magic are indeterministic, the answer is indeterminate: the probability that Merlin's spell was the causally effective one is 0.9, and the probability that Morgana's spell was the causally effective one is 0.1.

The crucial difference from the previous example is that, because the laws of witchcraft are truly indeterministic, there is no underlying mechanism that could be used as evidence that Merlin's spell, rather than Morgana's, was decisive for turning the prince into a frog (or vice versa). To see this, consider a variation of the scenario in which Merlin first casts a spell to turn the prince into a frog, and then Morgana casts a spell to turn him into a white rabbit. The two spells hang in the air until midnight, and then one of them takes effect, but which one depends on an indeterministic magical process. If the spells are different, one can look at the result after midnight and determine which one took effect. However, if they happen to be the same, this is impossible, and the probability that Merlin (whose spell came first) was the culprit is 0.9.

One might object that while this example may be more convincing than the previous one, since there is no underlying mechanism that could explain away the indeterminacy, it introduces a scenario that cannot be realized in the actual world. However, many scenarios discussed in the literature on causation are never realized in the actual world. This is true not only of Schaffer's magic-world scenarios, but of almost the entire debate about overdetermination, preemption, and trumping. If these scenarios are taken seriously because they tell us something relevant about the concept of causation, then indeterministic trumping scenarios should also be taken seriously because they tell us something relevant about the concept of causation, namely that causal claims can be indeterminate.

A more general objection to the cases presented in this section is that assigning probabilities to entire causal claims contradicts our ordinary understanding of the relationship between causation and probability. We can assign probabilities to events, such as the event of Suzy's kicking a ball or Morgana's casting a spell. In addition, we can ascribe conditional probabilities to events given that certain other events occur. For instance, we can understand the probability that the ball would have hit the amphora if Jimmy had not caught it as the conditional probability of the ball's hitting the amphora given that Jimmy does not catch it. But even though several theories of causation rely on such ascriptions of (conditional) probabilities to events (for example, Spirtes, Glymour, and Scheines 2000; Woodward 2003), they usually do not ascribe probabilities to the causal claims themselves.

However, the observation that something is unusual or has not been considered as an option before, does not imply that it does not make sense. Moreover, as emphasized above, the indeterministic cases discussed in this section are only one way to argue for causal indeterminacy. In particular, the Sorites paradoxes discussed in the previous section cannot be analyzed by assigning probabilities to statements describing borderline cases. For instance, if the probability that the irrigation system failed is 0.6, this does not imply that the probability that the failure of the system caused the death of the plants is also 0.6 (or has any other specific value).

And the same holds mutative mutandis for the headache pill case.⁷ This underscores the general thesis of the paper: causal indeterminacy occurs in scenarios with quite different causal structures, and even if one rejects one or two of them, there are still others that support it. In the next and last section, I discuss two possible general objections to my argument.

4. Two Possible Objections

One possible reaction to the conclusion that there is indeterminacy in causation is to argue that the supposedly indeterminate causal statements have not been understood in the right way. Several authors have pointed out that causation should be understood as a contrastive relation rather than as a relation between two individual events (Schaffer 2005; see also Hitchcock 1996; Woodward 2003: 67–68; Maslen 2004). So, would the phenomenon of causal indeterminacy disappear if all causal statements were reformulated in such a way that the relevant contrasts on both the cause and the effect side were made explicit?

Consider the janitor example from the beginning. There is a non-contrastive causal statement about this scenario, which does not have a determinate truth value:

- (*) The janitor's doing nothing was a cause of Jane having to cancel her class.

There is a contrastive reformulation of (*) which is determinately true:

- (**) The janitor's doing nothing rather than fixing Jane's classroom caused Jane having to cancel her class rather than being able to teach it.

However, there is also a contrastive reformulation of (*) which is indeterminate:

- (***) The janitor's doing nothing rather than fixing *one of the classrooms* was a cause of Jane having to cancel her class rather than being able to teach it.

Obviously, the conclusion that there is indeterminacy in causation should not depend upon the implausible claim that *all* contrastive causal statements that can be made about a certain situation are indeterminate (see also fn. 2). Moreover, there are usually several possibilities to turn a non-contrastive causal statement into a contrastive one, and none of these possibilities is privileged over the others: they just express different contrastive causal claims. Accordingly, if one is a contrastivist about causation and if *there are* indeterminate contrastive causal claims, this will be sufficient for concluding that there is indeterminacy in causation.⁸

⁷ According to some approaches to vagueness, borderline cases are analyzed in terms of degrees of truth, which have formal similarities with probabilities. However, it is widely agreed that degrees of truth should not be equated with probabilities (for discussion see Keefe 2000: ch. 4).

⁸ Wilson proposes a determinable-based account of metaphysical indeterminacy, according to which metaphysical indeterminacy 'involves an object's (i) having a determinable property, but (ii) not having any unique determinate of that determinable' (Wilson 2013: 359). A central challenge for this approach is to make sense of the idea that whole causal claims can be determinates of other causal claims, for instance, that the claim 'the janitor's failure to fix Jane's classroom caused Jane to cancel her class' is, in a sense to be specified, a determinate of the determinable 'the janitor's failure to fix any of the classrooms caused Jane to cancel her class'. A discussion of this question is beyond the scope of this paper.

In addition, there are indeterminate contrastive statements about each of the other scenarios presented. In the headache pill example (Section 2), it is plausible to assume that the contrastive statement that Jimmy's taking the second pill rather than not taking it caused his headache to disappear rather than linger has the same truth value as the non-contrastive statement that Jimmy's taking the second pill caused his headache to disappear. But then, the contrastive statement is indeterminate iff the non-contrastive statement is. And the same holds, *mutatis mutandis*, for the other scenarios. Therefore, the conclusion that there is indeterminacy in causation cannot be avoided by pointing out that causation should be understood as a contrastive relation.

A second objection to the argument of this paper is that each of the scenarios used to show that there is causal indeterminacy relies on certain metaphysical assumptions. The janitor and irrigation system examples do not work if there is no absence causation. The headache pill example is problematic if events always have precise temporal boundaries. And the three cases presented in the previous section assume that there can be truly indeterministic processes. Should one not try to avoid the conclusion of this paper by rejecting all these presuppositions?⁹

First, note that the argument follows a common methodology in debates about the general features of the causal relation, for example, in the debates about whether causation is a transitive relation, whether causal statements are contrastive (as mentioned above), or whether causes always precede their effects in time. These debates are typically based, at least in part, on intuitive judgments about fictional scenarios that rely on specific metaphysical assumptions. For example, arguments that causation is not a transitive relation (Hitchcock 2001) are based on scenarios involving higher-level events. They thus rely on the assumption that there can be higher-level causation and that causal relations do not occur only at the physical level. Schaffer's argument that causation is a contrastive relation is partly driven by the consideration that there should be an illuminating approach to absence causation (Schaffer 2005: 329–32). Therefore, the mere fact that the scenarios rely on certain metaphysical assumptions, does not mean that they do not have any general implications for our understanding of causation.

The second thing to note is that there is not just one specific metaphysical assumption that triggers the intuition that there may be indeterminacy in causation. The examples presented not only have quite different causal structures but are also based on different metaphysical assumptions. Each of these assumptions can be defended independently of the others and independently of the discussion of causal indeterminacy. This should be sufficient evidence that causal indeterminacy should be taken seriously.

A final note: the conclusion that some causal statements are indeterminate is primarily a claim about the general features of the *concept* of causation, and this raises the further question, mentioned at the beginning, of whether causal indeterminacy is a kind of worldly or ontological indeterminacy. This question deserves much

⁹ Yet another strategy to call the argument of this paper into question consists in stipulating that there are two different notions of causation, one of which gives rise to indeterminacy, whereas the other one does not. A prominent proponent of the view that there is not just one unified concept of causation is Hall (2004); a related view, relying on the distinction between 'is a cause of' and 'is causally relevant to', has been defended by Swanson (2010). However, this strategy appears promising only if the cases leading to causal indeterminacy share a characteristic that systematically distinguishes them from other cases, and this is unlikely given the diversity of the cases leading to causal indeterminacy.

more attention than has been given here—but to discuss it would go beyond the scope of the present argument.

Acknowledgements

This paper has a rather long history and has benefited from suggestions and comments from many people and audiences. An earlier version was presented at the 2018 conference of the Society for the Metaphysics of Science. Max Kistler commented on the paper at that conference and gave me incredibly detailed and constructive feedback, which helped me to clarify my thinking on the topic. Munjong Kolss, Nan Li, Vivek Mathew, Matthias Rolffs, and Jonas Werner read earlier versions of the manuscript and provided extremely valuable input in many ways. I have also substantially benefited from discussions with Sara Bernstein, Andreas Hüttemann, the members of the Rutgers Metaphysics Group in spring 2017, and the members of the eidos Center for Metaphysics in Geneva. Finally, I would like to thank Antony Eagle and two anonymous reviewers for the *Australasian Journal of Philosophy* for providing extremely detailed and helpful comments on earlier versions of the manuscript.

Disclosure Statement

No potential conflict of interest was reported by the author(s).

Funding

This research was supported by the German Academic Foundation (Deutsche Forschungsgemeinschaft (DFG)), grants no. HO 5206/2-1 and HO 5206/3-1, and by the Swiss National Science Foundation (Schweizerischer Nationalfonds (SNF)), grant no. 100019E_189589.

ORCID

Vera Hoffmann-Kolss  <http://orcid.org/0000-0001-5324-0559>

References

- Ballarín, R (2014) 'Disjunctive Effects and the Logic of Causation', *British Journal for the Philosophy of Science* **65**: 21–38. doi:10.1093/bjps/axs040.
- Barnes, E (2010) 'Arguments Against Metaphysical Indeterminacy and Vagueness' *Philosophy Compass* **5**: 953–64. doi:10.1111/j.1747-9991.2010.00348.x.
- Beebe, H (2004) 'Causing and Nothingness', in J Collins, E J Hall, and L A Paul, eds., *Causation and Counterfactuals*: 291–308. MIT Press.
- Bernstein, S (2014) 'Omissions as Possibilities', *Philosophical Studies* **167**: 1–23. doi:10.1007/s11098-013-0229-0.
- Bernstein, S (2016a) 'Causal and Moral Indeterminacy', *Ratio* **29**: 434–47. doi:10.1111/rati.12144.
- Bernstein, S (2016b) 'Omission Impossible', *Philosophical Studies* **173**: 2575–89. doi:10.1007/s11098-016-0672-9.
- Bernstein, S (2017) 'Causal Proportions and Moral Responsibility', in D Shoemaker, ed., *Oxford Studies in Agency and Responsibility*, Vol 4: 165–82. Oxford University Press.
- Bunzl, M (1979) 'Causal Overdetermination', *Journal of Philosophy* **76**: 134–50. doi:10.2307/2025525.
- Collins, J (2000) 'Preemptive Prevention', *Journal of Philosophy* **97**: 223–34. doi:10.2307/2678391.
- Demirtas, H (2022) 'Causation Comes in Degrees', *Synthese* **200**: 1–17. doi:10.1007/s11229-022-03507-2.
- Dowe, P (2000) *Physical Causation*. Cambridge University Press.
- Ellis, R D (2004) 'Three Arguments against Causal Indeterminacy', *Philosophia* **31**: 331–44. doi:10.1007/bf02385189.

- Hall, N (2004) 'Two Concepts of Causation', in J Collins, N Hall, and L A Paul, eds., *Causation and Counterfactuals*: 225–76. MIT Press.
- Hitchcock, C (1996) 'The Role of Contrast in Causal and Explanatory Claims', *Synthese* **107**: 395–419. doi:10.1007/bf00413843.
- Hitchcock, C (2001) 'The Intransitivity of Causation Revealed in Equations and Graphs', *Journal of Philosophy* **98**: 273–99. doi:10.2307/2678432.
- Jeffrey, R C (1992) 'Mises' Redux', in *Probability and the Art of Judgment*: 192–202. Cambridge University Press.
- Kaiserman, A (2018) "'More of a Cause": Recent Work on Degrees of Causation and Responsibility', *Philosophy Compass* **13**. doi:10.1111/phc3.12498.
- Kaiserman, A (2021) 'Responsibility and the "Pie Fallacy"', *Philosophical Studies* **178**: 3597–616. doi:10.1007/s11098-021-01616-1.
- Keefe, R (2000) *Theories of Vagueness*. Cambridge University Press.
- Lewis, D (1973) 'Causation', *Journal of Philosophy* **70**: 556–67. doi:10.2307/2025310.
- Maslen, C (2004) 'Causes, Contrasts, and the Nontransitivity of Causation', in J Collins, N Hall, and L A Paul, eds., *Causation and Counterfactuals*: 241–57. MIT Press.
- McDermott, M (1995) 'Redundant Causation', *British Journal for the Philosophy of Science* **46**: 523–44. doi:10.1093/bjps/46.4.523.
- Sartorio, C (2004) 'How to Be Responsible for Something without Causing It', *Philosophical Perspectives* **18**: 315–36. doi:10.1111/j.1520-8583.2004.00031.x.
- Sartorio, C (2006) 'Failures to Act and Failures of Additivity', *Philosophical Perspectives* **20**: 373–85. doi:10.1111/j.1520-8583.2006.00111.x.
- Schaffer, J (2000) 'Trumping Preemption', *Journal of Philosophy* **97**: 165–81. doi:10.2307/2678388.
- Schaffer, J (2004) 'Causes Need Not Be Physically Connected to Their Effects: The Case for Negative Causation', in C R Hitchcock, ed., *Contemporary Debates in Philosophy of Science*: 197–216. Blackwell.
- Schaffer, J (2005) 'Contrastive Causation', *Philosophical Review* **114**: 327–58. doi:10.1215/00318108-114-3-327.
- Spirtes, P, C Glymour, and R Scheines (2000) *Causation, Prediction and Search*, 2nd ed. MIT Press.
- Swanson, E (2010) 'Lessons From the Context Sensitivity of Causal Talk', *Journal of Philosophy* **107**: 221–42. doi:10.5840/jphil2010107517.
- Swanson, E (2017) 'Indeterminacy in Causation', *Philosophical Quarterly* **67**: 606–24. doi:10.1093/pq/pqw068.
- Torrigo, L (2000) 'Vague Causation', *Nous* **34**: 313–47. doi:10.1111/0029-4624.00214.
- Wilson, J M (2013) 'A Determinable-based Account of Metaphysical Indeterminacy', *Inquiry* **56**: 359–85. doi:10.1080/0020174x.2013.816251.
- Woodward, J (2003) *Making Things Happen: A Theory of Causal Explanation*. Oxford University Press.