Three Concepts of Actual Causation

Enno Fischer

Abstract

I argue that we need to distinguish between three concepts of actual causation: total, path-changing, and contributing actual causation. I provide two lines of argument in support of this account. First, I address three thought experiments that have been troublesome for unified accounts of actual causation, and I show that my account provides a better explanation of corresponding causal intuitions. Second, I provide a functional argument: if we assume that a key purpose of causal concepts is to guide agency, we are better off making a distinction between three concepts of actual causation.

1 Introduction

At least since Lewis's ([1973]) pioneering work, questions of redundancy have been at the heart of discussions about actual causation. A straightforward counterfactual analysis would state that an event c_1 is an actual cause of another event e if and only if it is true that e would not have occurred without c_1. But this approach fails when there are alternative events $c_2, c_3, ...$ that would have brought about e in the absence of c_1. A standard example are instances of pre-emption. Suppose that an assassin in training shoots a victim. Suppose also that the trainee is

accompanied by a supervising assassin who would have shot the victim if the trainee had failed to do so. We clearly identify the trainee as actual cause in the given situation even though the victim's death does not depend on the trainee's actions (Hitchcock [2001]).

The discussion about actual causation and redundancy has been considerably advanced through the use of causal models as developed by Spirtes *et al.* ([1993]) and Pearl ([2000]), among others. The framework is useful because it enables us to represent complex relations of counterfactual dependence. There is now a broad consensus that causal models are also a powerful tool for representing relations of actual causation. Yet there is no consensus on how exactly concepts of actual causation are to be defined within this framework.

In the causal modelling literature there are, broadly speaking, two approaches to defining concepts of actual causation. Both consist in weakening the original idea of direct counterfactual dependence, but they disagree about how this has to be done. First, there is a class of accounts that require counterfactual dependence while certain factors other than the cause may be held fixed in their actual state (Pearl [2000]; Hitchcock [2001]; Halpern [2016]). For example, keeping fixed the actual fact that the supervisor does not shoot (no matter whether the trainee fails or not), there is a counterfactual dependence of the victim's death on the trainee's actions. This is taken to be why the trainee is an actual cause. Second, there is a class of accounts that are more permissive in that they allow other factors to be set to non-actual states (Pearl [2000]; Halpern and Pearl [2005]; Halpern [2016]). These accounts have generally been considered to be more powerful than the accounts in the first class because they also capture our causal intuitions as evoked by scenarios involving symmetrical overdetermination. At the same time, they have been criticized as being too permissive (Hall [2007]; Livengood [2013]; Weslake [unpublished]).

A widely held assumption in this debate is that there is a single concept of actual causation. This concept is taken to be captured either by a definition from the first class or by a definition from the second class (or some entirely different account that is yet to be developed). In this article I will argue that this assumption is misguided. I instead contend that there are at least three concepts of actual causation that stand in a hierarchical relation: total, path-changing, and contributing actual causation. The concept of total actual causation (TAC) amounts to

straightforward counterfactual dependence. This is the simplest concept of actual causation and it applies to scenarios that do not involve redundancy. If c is a TAC of e, then it will be possible to prevent e by intervening on c (and no other factor). The concept of path-changing actual cause (PAC) is more permissive than the concept of TAC in the sense that it covers all instances of TAC but also those instances where a counterfactual dependence between two factors c and e is revealed only if certain other factors are held fixed in their actual state. If c is a path-changing actual cause of e, but not also a TAC, then it will not be sufficient to intervene on c in order to prevent e. Instead one has to combine the primary intervention on c with secondary interventions that counteract the adverse consequences of the primary intervention. Finally, the concept of contributing actual cause (CAC) is the widest concept of actual causation in the sense that it covers all instances of TAC and PAC but also those where a counterfactual dependence between c and e is only revealed if certain other factors are set to non-actual values. This means that in order to prevent an effect e that is caused by several contributing actual causes c_i, each of the contributing causes has to be targeted with an independent intervention.

I provide two lines of argument in support of the distinction between total, path-changing, and contributing actual causation. First, I show that distinguishing between these three concepts helps account for three example cases that have been troublesome for extant unified accounts: trumping (Schaffer [2000]), the light bulb (Halpern [2016]), also known as 'Shock C' (McDermott [1995]), and the henchman (Rosenberg and Glymour [2017]). In each of these examples, there are two factors that both seem to qualify as actual causes. Yet, there is an important asymmetry between these factors. Extant theories face a dilemma here: either they describe both factors as actual causes (and thus cannot explain the asymmetry), or they dismiss the intuition that both factors are actual causes (and account for the asymmetry by identifying only one of the factors as actual cause). By distinguishing between different kinds of actual causes, my account can both hold that each example involves two actual causes, and explain the perceived asymmetry.

The second line of argument is related to considerations about the function of causal concepts. I assume that an important function of such concepts is to indicate suitable targets of intervention: if event c is an actual cause of event e, then preventing c is a suitable strategy

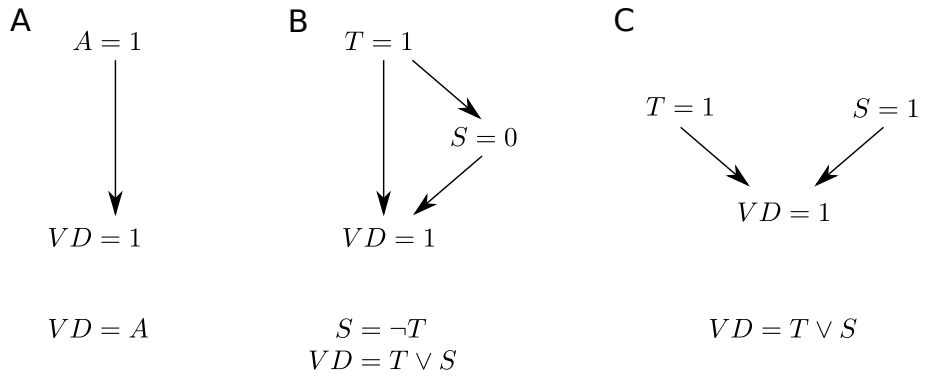

Figure 1. Three examples of actual causation. (A) Straightforward dependence. (B) Pre-emption. (C) Symmetrical overdetermination.

for preventing e. But in order to be informative, it should be at least clear whether such an intervention is sufficient. In cases of total actual causation a single intervention is sufficient, in cases of path-changing and contributing actual causation this is not the case.

The article is structured as follows: In Section 2 I will briefly review existing accounts of actual causation that have been proposed within the framework of causal models, and introduce the concepts of total, path-changing, and contributing actual causation. According to an underlying assumption in this debate, only one of these three concepts—or another, yet to be discovered concept—may truly qualify as 'the' concept of actual causation. In Section 3 I will challenge this view and argue that the definitions of total, path-changing, and contributing actual causation all describe legitimate causal concepts. In Section 4 I will present three example cases that raise problems for unified accounts and show that my approach provides a better account of these examples. In Section 5 I will provide some clarifications and relate my account to existing accounts of actual causation. In Section 6 I will present the functional argument. Section 7 provides a conclusion.

2 Total, Path-changing, and Contributing Actual Causation

The concepts of actual causation will be defined in the formal framework of causal models. A causal model M is an ordered pair, $\langle \mathcal{V}, \mathcal{E} \rangle$, where $\mathcal{V}$ is a set of variables and $\mathcal{E}$ a set of structural equations. The variables in $\mathcal{V}$ are represented by capital letters $(X, Y, \dots)$ and they have two or more possible values, represented by lower case letters $(x, y, \dots)$. These values represent potential events in the model's target system. A variable's taking on one of the values

(for instance $X = x$) represents an actual event in the target system. The structural equations $\mathcal{E}$ provide a summary of the counterfactual dependencies that hold between the events. As an example consider the scenario where an assassin pulls the trigger of her gun ($A = 1$) and kills a victim ($VD = 1$). The counterfactual dependence is then expressed by the structural equation $VD = A$: if the assassin had not pulled the trigger ($A = 0$), then the victim would not have died ($VD = 0$). The corresponding representation in terms of a directed acyclic graph is displayed in Figure 1A.

Here is a straightforward attempt to define actual causation within this framework:

Total Actual Cause: $X = x$ is a total actual cause of $Y = y$ in $(M, \vec{u})$ if and only if

TAC1: $(M, \vec{u}) \models (X = x) \wedge (Y = y)$.

TAC2: There exists a setting x' of variable X such that

$$(M, \vec{u}) \models [X \leftarrow x']\neg(Y = y).$$

Condition TAC1 states that in order for an event $X = x$ to be an actual cause of another event $Y = y$ it needs to be the case that in the current situation both $X = x$ and $Y = y$ occur. Condition TAC2 requires that if event $X = x$ had not occurred, then also the effect ($Y = y$) would not have occurred. With regard to the assassination example this definition seems perfectly adequate. The assassin's pulling the trigger of her gun ($A = 1$) satisfies both these conditions and we identify her as actual cause of the victim's death.

Total actual causation (TAC) represents an understanding of actual causation that is widely employed, for example, in the law, where it features as the 'but for' criterion of actual causation. However, as a definition of actual causation it is typically rejected because it does not apply to circumstances that involve redundancy. We will look at two kinds of redundancy. First, there are cases of pre-emption. Suppose that this time the victim is shot by an assassin in training T who is accompanied by a supervising assassin S. If the trainee had lost her nerve, then the supervisor would have pulled the trigger of her gun and the victim would still have died (see Figure 1B for the causal model). The problem is: we still identify the trainee as actual cause—even though the victim's death does not depend counterfactually upon her actions.

However, if we keep fixed the fact that the supervisor does not pull the trigger of her gun (as is the case in the actual situation), then it is true that the victim's death depends upon the trainee's actions. More generally, counterfactual dependence can sometimes be recovered by considering that factors other than the cause remain in their actual state. Accounts along these lines have been proposed by Pearl ([2000]), Hitchcock ([2001]), and Halpern ([2016]). The key idea of these accounts is captured by the following definition of path-changing actual cause (PAC):

> Path-changing Actual Cause: $X = x$ is a path-changing actual cause of $Y = y$ in $(M, \vec{u})$ if and only if
>
> PAC1: $(M, \vec{u}) \models (X = x) \wedge (Y = y)$.
>
> PAC2: There is a set $\vec{W}$ of variables in $\mathcal{V}$ and a setting x' of variable X such that if $(M, \vec{u}) \models \vec{W} = \vec{w}^*$, then
>
> $$(M, \vec{u}) \models [X \leftarrow x', \vec{W} \leftarrow \vec{w}^*]\neg(Y = y).$$

PAC1 is equivalent to TAC1 and states that both the cause and the effect have to occur. PAC2 is a slightly more sophisticated counterfactual conditional. Like TAC2 it requires a possible setting of the cause variable $X = x'$ that changes the value of effect variable Y. But this time the change in the effect variable may be achieved while other variables $\vec{W}$ are held fixed at their actual value $\vec{w}^*$. Here the vector notation indicates that there may be more than one variable in the set of variables $\vec{W}$, and $\vec{w}$ assigns a value to each of these variables. Note that path-changing actual causation is a strictly more encompassing concept than total actual causation. More specifically, total actual causation describes the special case where $\vec{W}$ is the empty set.

I choose the label 'path-changing actual cause' because intervening on such a cause changes the causal path along which the effect is influenced. If I intervene on the trainee such that she does not pull the trigger, I will not prevent the victim's death. But as a result of the intervention it will be the supervisor's bullet that kills the victim. That is, the intervention changes along which of the two causal paths linking the variables T and VD in Figure 1B the causal process is running.

The second kind of redundancy is called symmetrical overdetermination and involves two factors that both bring about the effect at the same time, while either one of the factors would have been sufficient. Suppose, for example, that this time both the trainee and the supervisor shoot at the same time and hit the victim at the same time such that each one of the two bullets would have been lethal (see Figure 1C). We typically identify both the trainee and the supervisor as actual causes even though the victim's death depends on neither of the individual agents' actions. This scenario is typically taken as a reason for rejecting definitions along the lines of path-changing actual causation.[1] The reason is that even if we hold fixed the fact that the supervisor shoots (as she does in the actual situation), there is no way to restore the counterfactual dependence of the victim's death on the trainee's actions. Instead we need to set the supervisor variable to a non-actual value. This idea is captured by the definition of contributing actual cause (CAC).

> Contributing Actual Cause: $X = x$ is a contributing actual cause of $Y = y$ in $(M, \vec{u})$ if and only if
>
> > CAC1: $(M, \vec{u}) \models (X = x) \wedge (Y = y)$.
> >
> > CAC2: There exists a partition $(\vec{Z}, \vec{W})$ of $\mathcal{V}$ with $X \subseteq \vec{Z}$ and some setting $(x', \vec{w}')$ of the variables in $(X, \vec{W})$ such that if $(M, \vec{u}) \models Z = z^*$ for all $Z \in \vec{Z}$, then both of the following conditions hold:
> >
> > (a) $(M, \vec{u}) \models [X \leftarrow x', \vec{W} \leftarrow \vec{w}']\neg(Y = y)$.
> >
> > (b) $(M, \vec{u}) \models [X \leftarrow x, \vec{W}' \leftarrow \vec{w}', \vec{Z}' \leftarrow \vec{z}^*](Y = y)$ for all subsets $\vec{W}'$ of $\vec{W}$ and all subsets $\vec{Z}'$ of $\vec{Z}$.

The definition of contributing actual cause is closely related to definitions suggested by Pearl and Halpern (Pearl [2000]; Halpern and Pearl [2005]). This definition is even more permissive than the definition of PAC because CAC2(a) allows us to restore counterfactual dependence between cause and effect by setting other variables also to non-actual values $\vec{w}'$. However,

[1] An exception is Halpern ([2016]) who takes his 'modified' definition of actual causation to account for instances of symmetrical overdetermination as well. This is achieved by assuming that the conjunction of the individual causes stands in a relation of counterfactual dependence with the effect. For reasons detailed in Section 6 I do not take this to be a legitimate treatment of symmetrical overdetermination.

CAC2(a) alone would be too permissive. Going back to the pre-emption example, suppose we add another variable $SR = 1$ feeding into variable S such that $S = SR \wedge \neg T$. This variable may represent that the supervisor is ready to act in case the trainee fails to pull the trigger. Condition CAC2(a) would falsely identify $SR = 1$ as actual cause. This is why we need the additional requirement detailed in CAC2(b), which ensures that the change in Y is attributed to an actual cause. With this, I refer interested readers to Halpern and Pearl's ([2005], pp. 852–59) detailed treatment of this condition; these details will not be relevant to the subsequent discussion.

Note that the concept of CAC is strictly more encompassing than the concept of PAC. Every PAC is also a CAC, but not vice versa. More specifically, PAC describes the special case where the values of variables in $\vec{W}$ may only be set to values $\vec{w}^*$ that these variables take on in the actual situation.

The definition of CAC is so powerful since it yields the correct verdict in scenarios that involve pre-emption and symmetric overdetermination—two kinds of scenarios that have been notoriously difficult to handle within counterfactual accounts. This is achieved by loosening the requirement of counterfactual dependence making it a more permissive concept than the concept of total and path-changing actual causation. However, as we will see in Section 4, there are also a number of examples that have been taken to show that the concept of CAC is too permissive.

3 A Unified Account of Actual Causation?

A key assumption underlying the foregoing discussion is that there is a unified concept of actual causation. Contributions to the literature typically provide counterexamples to some existing definition of actual causation and offer a new definition that covers the counterexamples as well as those examples that were also covered by the old definition. The new definition is considered to be a better approximation of an underlying concept of actual causation and is thought to replace earlier suggestions.

The assumption of unification is particularly salient in (Halpern [2016]). Halpern offers three competing basic definitions of actual causation: the 'original' and the 'updated' HP definitions (which are similar to the definition of CAC), and the 'modified' HP definition (which is similar

to the definition of PAC), and discusses several ways of implementing considerations related to default and deviant values. At the same time, Halpern seems to be convinced that there is one unified concept of actual causation that captures the causal intuitions in the wealth of example cases and he suggests that his 'modified' HP definition comes closest.

The assumption of unification is endorsed even by those who—on another level—accept that there is a plurality of causal concepts. Woodward ([2003]) provides a detailed analysis of the relation between causation and manipulability. The result is a theory that distinguishes between total causes, direct causes, contributing causes, and actual causes. Thus, while Woodward accepts that there is a plurality of causal concepts, he seems to assume that, within that plurality, 'actual causation' refers to a unified concept.

In contrast, I suggest that the definitions of TAC, PAC, and CAC all describe legitimate concepts of actual causation. That is, I argue that there is not one unified concept of actual causation but that there are at least three different concepts of actual causation. As a consequence, it will occur that there is a sense in which some $X = x$ is an actual cause of $Y = y$ while there is another sense in which it is not true that $X = x$ is an actual cause of $Y = y$. For example, it is true that the trainee is a CAC of the victim's death in the scenario involving symmetrical overdetermination. But it is not true that the trainee is a TAC.

What do we gain by describing TAC, PAC, and CAC as legitimate concepts? In the following section I will show that this helps to clarify three example cases that have been troublesome for extant unified accounts. In each of these cases one faces difficulties if one takes only one of the extant concepts of actual causation to be legitimate. But if we accept that TAC, PAC, and CAC are all legitimate concepts, then the causal intuitions evoked by the examples can be explained. After some clarifications in Section 5 I will argue in Section 6 that by identifying TAC, PAC, and CAC as legitimate concepts we can also provide a better account of how causal claims guide agency.

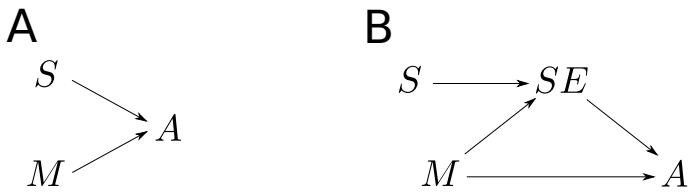

Figure 2. Two models for trumping.

4 Three Problem Cases

4.1 Trumping

A major and a sergeant order a corporal to advance. The corporal receives both orders at the same time and advances. What caused the corporal to advance? The case is ambiguous. On the one hand, the corporal receives both orders at the same time. Thus, it seems like both orders are causes just as in situations of symmetrical overdetermination. On the other hand, orders of higher-ranked officers trump those of lower-ranked officers. Thus, it seems like there is an important asymmetry between the two orders such that some have argued that the only cause is the major's order (Schaffer [2000]; Lewis [2004]).

Let us consider treatments of this case in terms of contributing actual causation because these seem to correspond to the most advanced accounts, such as the one proposed by Halpern and Pearl ([2005]). Halpern and Pearl are aware of the ambiguity of trumping cases and they offer two possible models. The first is a coarse-grained model that consists of three variables (see Figure 2A): there are two variables, M and S, representing the major's and the sergeant's actions, respectively. These have three possible values: 1 (order advance); -1 (order retreat); 0 (do nothing). Moreover, variable A represents the corporals actions. If the major issues an order, the corporal follows it, that is, $A = M$ if $M \neq 0$. If the major does not issue an order, the corporal follows the sergeant's orders, that is, $A = S$ if $M = 0$. In the actual situation both the major and the sergeant order to advance and the corporal advances: $M = S = A = 1$. With this model the definition of contributing actual causation identifies both the major and the sergeant as actual causes of the corporal's action.

But what about the perceived asymmetry between the major and the sergeant? Halpern and Pearl ([2005]) argue that there is an alternative model according to which only the major is an actual cause. The alternative model includes a variable SE that captures whether the sergeant's

order is effective. If the major does nothing ($M = 0$), then $SE = S$. But if the major issues an order ($M \neq 0$), then the sergeant's order is not effective ($SE = 0$; see Figure 2B).

This model reproduces Schaffer's and Lewis's claim that only the major's issuing the order is an actual cause. However, this does not seem to be a legitimate treatment of trumping cases, at least if they are construed as cases where both causal processes run to completion.

But even if this treatment were legitimate, Halpern and Pearl's account would not be completely satisfactory. According to the coarse-grained model, both the major and the sergeant are actual causes and, according to the fine-grained model, only the major is an actual cause. But what are the reasons for preferring one model over the other? Moreover, according to these models, it is either the case that the major and the sergeant are to be treated on a par or it is the case that only one of them is an actual cause. But the intuition seems to be that both officers are actual causes and that there is an asymmetry between them at the same time.

A better treatment of trumping cases is available if we acknowledge that there are other legitimate concepts of actual causation besides the concept of contributing actual causation. Consider once more Halpern and Pearl's simple model of the case. According to this model, both the sergeant and the major are identified as contributing actual causes. Given that the major does not issue an order, the corporal's behaviour depends upon the sergeant's order. The same is true for the major. Given that the sergeant does not issue an order, the corporal's actions depend on the major's actions.

However, there are also actions available to the major such that the corporal's behaviour depends on the major regardless of the sergeant's actions. More specifically, the major could have ordered the corporal to retreat. Since the corporal then would have retreated (no matter what the sergeant orders), it is clear that the major is also a total actual cause of the corporal's behaviour. Thus, we can accommodate the intuition that both the major and the sergeant are actual causes. At the same time, the perceived asymmetry is explained by the fact that only the major is a total actual cause.

Halpern and Pearl appear to endorse a similar account of the asymmetry.[2] They point out that the major is a strong cause while the sergeant isn't ([2005], p. 874). The definition of strong

[2] Thanks to an anonymous referee for bringing up this point and the following point about Lewis and Schaffer.

cause builds upon the definition of contributing actual cause (or actual cause in Halpern and Pearl's terminology) and involves an extra condition requiring that

$$(M, \vec{u}) \models [X \leftarrow x, \vec{W} \leftarrow \vec{w}''](Y = y)$$

for all settings $\vec{w}''$ of $\vec{W}$. The major's ordering to advance ($M = 1$) is a strong cause because for all possible values of the sergeant variable S the corporal advances if $M = 1$. By contrast, the sergeant is not a strong cause because the corporal does not advance if the major orders to retreat.

However, one difference between my account and Halpern and Pearl's account is that they reject the concept of strong cause where it does not coincide with their concept of actual cause: 'in many of our examples, causality and strong causality coincide. In the cases where they do not coincide, our intuitions suggest that strong causality is too strong a notion' ([2005], p. 855). Just as Halpern and Pearl's concepts of strong cause and actual cause my concepts of TAC and CAC coincide in most situations—situations that involve straightforward counterfactual dependence of the effect on the cause. Moreover, in instances where CAC and TAC do not coincide, the concept of TAC seems to be too restrictive to capture all those factors that are intuitively identified as causes. But I do not take this to be a reason to endorse the concept of CAC instead of the concept of TAC. Instead I argue that we need both the concept of TAC and the concept of CAC. Endorsing both these concepts is what enables us to accommodate the intuitions that apply to cases like trumping, where two factors are causes in the weaker sense of CAC and only one factor is a cause in the stronger sense of TAC.

There is a different question of whether the distinction between actual cause and strong cause can give rise to a viable pluralist account if one endorses both concepts, even in cases where the two concepts do not coincide. I think that such a pluralist account could in principle be developed. However, I prefer a pluralist account in terms of TAC, PAC, and CAC because it also accounts for the two examples to be discussed shortly, whereas the distinction between actual and strong causation does not help in those examples.

I have argued that my account explains a perceived asymmetry between the major and the sergeant. But presumably this will not satisfy authors like Lewis ([2004]) and Schaffer ([2004])

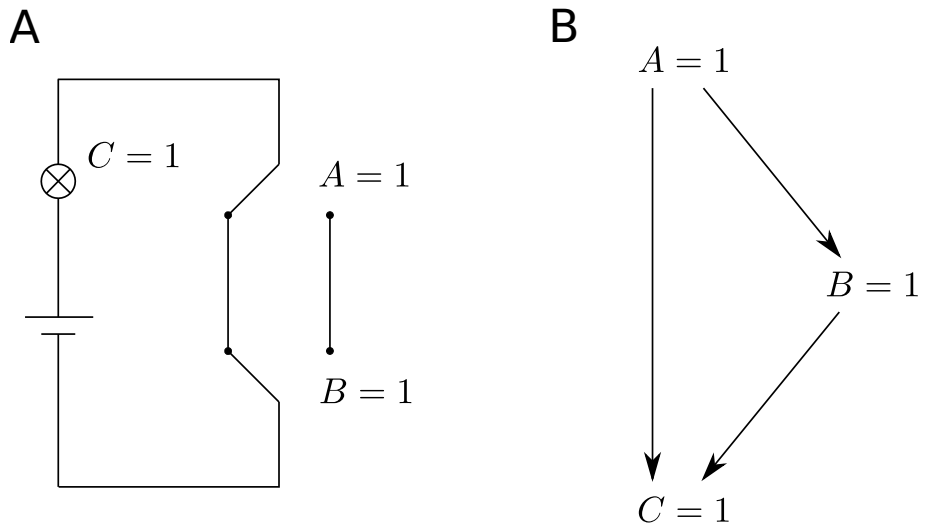

Figure 3. The light bulb case. If both switches are either connecting to the left or right wire, then the lamp C is switched on. (A) The circuit. (B) The causal graph.

who claim that the sergeant is not a cause at all. Note that Lewis and Schaffer do not just have a different way of accounting for certain basic causal intuitions. They also have a different view on what the basic causal intuitions are in the first place. But how do we know which are the 'correct' causal intuitions? In Section 6 I will address these concerns, arguing that identifying the sergeant as (contributing) actual cause makes sense from the perspective of an agent who aims to influence whether the corporal advances or not.

4.2 The light bulb

Here is another problematic case:

> A and B each control a switch. There are wires going from an electricity source to these switches and then continuing on to C. A must first decide whether to flip his switch left or right, then B must decide (knowing A's choice). The current flows, resulting in a bulb at C turning on, iff both switches are in the same position. B wants to turn on the bulb, so flips her switch to the same position as A does, and the bulb turns on. (Halpern [2016], p. 100)[3]

Figure 3A depicts a schematic representation of the kind of circuit that is described here. The structural equations of the case, following Halpern, are: $B = A$, saying that B copies the position of A's switch, and $C = 1$ if and only if $A = B$, saying that the light bulb is switched on if and only if the two switches are in the same position.

[3] This example is isomorphic to the scenario involving two agents A and B who use an electric circuit in order to apply shocks to agent C as discussed by McDermott ([1995]) and Weslake Weslake ([unpublished]).

Causal intuitions are again ambiguous. On the one hand, both agents' actions are causes of the fact that this particular circuit is closed. On the other hand, there is clearly an asymmetry between the two agents because only B has control over whether the light is switched on. This ambiguity is also reflected by Halpern's discussion. On the one hand, Halpern ([2016], p. 100) states that 'Intuition suggests that A's action should not be viewed as a cause of the C bulb being on, whereas B's should'. On the other hand, his treatment of the case in terms of graded causation suggests that he takes both agents to be causes, while the perceived asymmetry between them is a matter of degree and is to be explained in terms of considerations associated with normality ([2016], p. 102). The definition of contributing actual cause (as well as the definition of path-changing actual cause) treats both agents on a par and, thus, does not reflect the important fact that only B has control over the lamp's being switched on.

Acknowledging that there is not a single unified concept of actual causation, again, helps account for the example. B is a total actual cause of the light's being switched on because there is a relation of straightforward counterfactual dependence. A is also an actual cause, however, it is only a path-changing actual cause: we need to keep B fixed at its actual value, in order to reveal a counterfactual dependence of C on A.[4] Thus, we can account for the intuition that both A and B are actual causes of the light's being switched on but at the same time we can account for the intuition that B has better control over the lamp's being switched on.

4.3 The henchman

A gang leader ($GS = 1$) and his henchman ($HS = 1$) both shoot and their enemy dies ($ED = 1$). The enemy would also have died if either only the gang leader or only the henchman had shot. However, the henchman shoots if and only if the gang leader shoots. We shall also assume that the two individually lethal bullets kill instantaneously and hit at the same time, such that they stand in a relation of symmetrical overdetermination. The causal graph is displayed in Figure

[4] In order to reveal the path-changing character of A we need to choose a model that reflects whether the current runs along the left wire (LW) in Figure 3A or along the right wire (RW), for example, a model with structural equations as follows: $B = A$, $LW = A \wedge B$, $RW = \neg A \wedge \neg B$, and $C = LW \vee RW$.

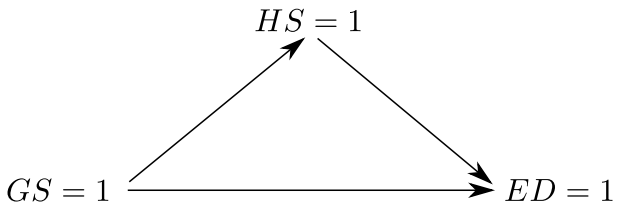

Figure 4. The henchman case. The henchman copies the gang leader's actions. The victim dies if either the gang leader or the henchman shoots.

4 and the structural equations are as follows:

$$HS = GS,$$

$$ED = GS \vee HS.$$

Clearly, both the gang leader and the henchman are to be identified as actual causes of the enemy's death. Yet, at the same time there seems to be an important asymmetry between the two agents because the enemy's death depends upon the gang leader's actions but not on the henchman's actions.

Rosenberg and Glymour ([2017]) present this case as a counterexample to Halpern's modified definition. According to Halpern's modified definition (and my definition of path-changing actual causation), the gang leader is an actual cause (because of the relation of counterfactual dependence) but the henchman is not an actual cause because there is no way to restore the counterfactual dependence of the enemy's death on the henchman's action by keeping the values of other variables fixed at their actual values.

The definition of contributing actual cause does identify both the gang leader and the henchman as actual causes. However, it treats both agents on a par and, thus, does not explain the intuitive difference between the two agents.

As in the other examples, a better explanation can be given if we accept that there is more than one legitimate concept of actual causation. The gang leader is a total actual cause of the enemy's death, meaning that he has complete control over the outcome of the situation. The henchman is also an actual cause but he is only a contributing actual cause: we need to set *GS* to a non-actual value in order to reveal the dependence of the enemy's death on the henchman's actions.

Let us take stock. I have presented three examples which all involve two factors that both seem to qualify as actual causes. Yet, there is also an important asymmetry between the two factors. Extant unified theories face a dilemma here: either they describe both factors as actual causes (and thus cannot explain the asymmetry), or they dismiss the intuition that both factors are actual causes (and account for the asymmetry by identifying only one of the factors as actual cause). By distinguishing between total, path-changing, and contributing actual causes, my account can both hold that each example involves two actual causes, and explain the perceived asymmetry.

5 Clarifications

My account involves three concepts of actual causation that all have close analogues in the existing literature. That is, I am not proposing a new definition of actual causation here. What is new, however, is that I suggest that the concepts of total, path-changing, and contributing actual causation are all legitimate concepts of actual causation. That is, I disagree with the commonly accepted assumption that the three concepts compete with each other.[5]

My approach has similarities with a number of pluralist proposals. Hall ([2004]), for example, argues that we need to distinguish between causation as dependence and causation as production. The dependence notion, according to Hall, accounts for the intuition that effects depend on their causes and can be applied to events as well as omissions. The dependence notion maps straightforwardly onto my concept of TAC. The production notion, according to Hall, accounts for the intuition that causation is transitive, local, and intrinsic. Hall argues that the production notion is in principle incompatible with the intuition that effects depend on their causes, especially in situations involving omissions. Even though the production notion seems to be the salient notion in examples involving redundancy such as those discussed above, I do not see a straightforward way to map it onto either one of my concepts of PAC or CAC.

It seems like Hall could accommodate the problem cases discussed in Section 4, arguing that

[5] A notable exception is Pearl's ([2000]) treatment. His approach includes a definition of actual cause (corresponding to my definition of PAC) from a definition of contributing cause (corresponding to my definition of CAC). However, Pearl does not provide a justification for distinguishing between these two concepts and the distinction is not retained in later publications, where a close analogue of his definition of contributing cause is taken to describe the concept of actual causation.

in each of the cases both factors are producers but only one of them involves dependence. But it is easy to change the examples such that they generate problems for Hall. One could, for example, make the henchman case a case involving omissions. Suppose a child is drowning in a lake and any attempt to save the child would be extremely risky because there is a violent thunderstorm. There is a chief lifeguard and an assistant lifeguard who can save the child only if they go out on the water together. The chief lifeguard is more experienced and her assistant always follows her decision. Because of the thunderstorm the chief lifeguard decides to stay on land, and so does the assistant. As a result the child drowns.

My theory says that both lifeguards are actual causes. The chief lifeguard is a TAC and the assistant is a CAC, analogous to the henchman case. According to Hall, the chief lifeguard is a cause in the sense of the dependence notion, but the assistant lifeguard is not a cause at all—which strikes me as unintuitive. The problem is that the lifeguards cannot be described as producers because the child's drowning results from omissions rather than actions.

A potential worry here is that of profligate causation (Menzies [2004], pp. 142–43). TAC essentially amounts to dependence. But if this includes dependence on omissions, then it seems that there is an endless number of causes of the child's drowning, including, for example, the fact that the Queen of England failed to save the child. Moreover, since the concept of CAC is even more permissive than the concept of TAC, the problem of profligate causation would be exacerbated.[6]

One way to deal with profligate causation is to impose suitable restrictions on $\mathcal{V}$, the set of variables that are part of the causal model. One could argue that an apt model of the situation should represent only serious possibilities (Halpern and Hitchcock [2010]).[7] The lifeguard's saving the child is such a serious possibility because it is her duty to save the child and she is sufficiently close. By contrast, the Queen's saving the child is not a serious possibility because we typically do not expect her to be around. Then the Queen's actions would not be an actual cause because the variable representing these actions would not be part of the causal model in the first place. A potential objection here is that imposing such restrictions on causal models

[6] Thanks to an anonymous referee for raising this issue.

[7] Alternatively, one can introduce a distinction between default and deviant values and add extra conditions to the definitions of TAC, PAC, and CAC that disallow counterfactual scenarios with too many deviant values, along the lines of (Halpern and Hitchcock [2015]).

makes the concepts of actual causation context sensitive and norm dependent. However, I do not think that this is a problem for my account. First, it reflects the fact that our judgements of actual causation are indeed often context sensitive and norm dependent (see empirical studies, such as Knobe and Fraser [2009]). Second, such context-sensitive and norm-dependent causal judgements are often particularly useful, for instance, when we deliberate about suitable strategies to achieve our goals (Hitchcock and Knobe [2009]; Hitchcock [2017]).

A further difference between Hall's and my approach is that my account is more optimistic with regard to causal models as a framework for defining concepts of actual causation. Hall thinks that a counterfactual analysis can be given only of his dependence notion of causation, while for the production concept we need a different kind of basic building block. Applying a terminology introduced by Hitchcock ([2007]) one could say that Hall proposes an extramural pluralism, meaning that he employs distinct kinds of basic building blocks (counterfactual conditionals, regularities) in order to generate a plurality of causal concepts. Then my account would be closer to an intramural pluralism because it generates a plurality of causal concepts employing only one sort of basic building block: (sophisticated) counterfactual conditionals framed in terms of causal models.

It will also be useful to compare my account with Cartwright's ([2007]) pluralist account. Cartwright argues that there is a wide variety of causal relations that is reflected by 'content-rich causal verbs', such as 'compress', 'attract', and 'discourage'. Cartwright ([2007], p. 21) states that on a unificationist view all these causal verbs are replaced by the abstract terms 'cause' and 'prevent' or, even worse, 'by one single piece of notation—the arrow [of a causal graph]'. This, she argues, is a problem because the content of the rich causal verbs is lost on such a view.

I agree with Cartwright that reducing all causal concepts to the kind of structural dependence that is expressed by the arrows of a causal graph is problematic. But I do not think that this is a problem for causal models. In fact, I take my account to illustrate how the framework of causal models can be employed to define a plurality of causal concepts. These causal concepts are surely not as specific as Cartwright's content-rich causal verbs. But neither are they as abstract as the arrows in a causal graph. Instead, they are located at an intermediate level of abstraction,

a level that I take to be abstract enough to be applicable to a wide range of circumstances and, at the same time, specific enough to give an agent clear guidance with regard to suitable targets of intervention (see the argument in Section 6).

There is another aspect of Cartwright's pluralist account, which is similar to the extramural pluralism of Hall. Cartwright argues that there are multiple theoretical frameworks that each account for certain paradigmatic features of causal reasoning but that each framework also has its limitations. She argues that the structural model framework, for example, faces problems where its central assumption of modularity does not apply (Cartwright [2007], p. 13). There are two ways the intramural pluralist can respond. The intramural pluralist can attempt to provide an account that exhausts all kinds of causal concepts such that extramural pluralism would not be needed. Alternatively, the internal pluralist can develop an intramural theory that tries to get as far as possible, but also acknowledge that there may remain instances of causation that will need other theoretical frameworks. On this view intramural and extramural pluralism coexist.

A conclusive response to Hall's and Cartwright's extramural pluralism is beyond the scope of my article because I focus on issues related to redundancy. But I take intramural pluralism to be able to make genuine progress even if it coexists with extramural pluralism. An advantage of my intramural account (over a purely extramural account) is that it clarifies the relationship between the different concepts of actual causation needed to account for cases involving redundancy. This is possible because the concepts are all defined in the same formal framework, that ultimately relates them to underlying relations of counterfactual dependence. The concept of CAC is the most permissive concept while the concepts of PAC and TAC describe special cases. More specifically, the three concepts exhibit a nested structure: TAC is a special case of PAC, and PAC is a special case CAC. This implies that every TAC is also a PAC and CAC, and that every PAC is a CAC—just as every square is a rectangle, and every rectangle is a quadrilateral (a plain figure with four edges and four vertices).

But doesn't the hierarchical relation between CAC, PAC, and TAC threaten the status of PAC and TAC as independent concepts of actual causation? That is, couldn't we state that the concept of CAC provides the basic analysis of what it means for an event to be an actual cause, and doesn't this imply that we have a unified concept of actual causation after all? It is true

that the concept of CAC is the most general concept. However, this does not make the concepts of PAC and CAC obsolete. I take the moral of the foregoing section to be that the concept of CAC (as well as the concept of PAC), if taken in isolation, is too coarse-grained. In order to capture the causal intuitions evoked by the example cases we need the concepts of CAC, PAC, and TAC.

6 A Functional Argument

So far my argument for distinguishing between total, path-changing, and contributing actual causation has been largely in the tradition of descriptive approaches to causation. Descriptive approaches appeal to common sense causal intuitions or causal judgements as the basic data and aim to provide definitions of causal concepts that account for this data. However, the descriptive project faces problems when there is disagreement about the basic data. As noted in Section 4, for example, Lewis and Schaffer do not share the intuition that in cases like the major-sergeant-corporal case both the major and the sergeant are actual causes. More generally, it seems a legitimate question to ask why we have certain causal intuitions and whether these intuitions are justified (or whether they simply reflect a biased view).

In this section I will turn to an evaluative stance. I will argue that we are justified in identifying in each of the three examples of Section 4 two actual causes and also justified in seeing an asymmetry between the respective causes. I shall argue that the concepts of TAC, PAC, and CAC are legitimate concepts because each one of them fulfils an important function. Following Woodward ([2014], pp. 693–94), the basic idea of such a claim about the function of causal concepts is to see causal cognition as an 'epistemic technology' and to evaluate it 'in terms of how well it serves our goals and purposes'. Following a major strand in the literature I will assume that causal reasoning is importantly linked to our ability to reach certain goals by means of manipulation (Cartwright [1979]; Woodward [2003]; Price [2017]). If I know that A stands in a causal relationship with B, I may be able to control B by manipulating A. This approach has been extended to claims of actual causation (Hitchcock and Knobe [2009]; Hitchcock [2017]; Danks [2013]). If I know that $A = a$ is an actual cause of $B = b$, this provides me with more specific information. I learn that I may prevent $B = b$ by intervening on A such that it takes on

a value $a' \neq a$.

My claim is that how exactly $B = b$ is to be prevented depends on the kind of actual cause that $A = a$ is. The difference between TAC, PAC, and CAC matters from the perspective of an intervening agent. The distinction is thus not merely a reaction to the problems that extant unified accounts of actual causation face. Instead it is positively justified by considerations concerning why we do and should have concepts of actual causation in the first place.

Suppose I know that the assassin in training went on her mission alone, without her supervisor (Figure 1A). If I want to save the victim, I have to intervene on the trainee. A simple intervention on the TAC is sufficient in order to accomplish my goal. Compare this with a situation of symmetrical overdetermination (where the trainee is a CAC). Here I have to intervene on all CACs. If we take concepts of actual causation to inform us about suitable targets of intervention, they should at least be able to reflect this important difference.[8]

In the assassin example there are only two contributing actual causes and both are straightforwardly identified as such. But the notion of contributing cause is more general. It also applies to situations where the effect is multiply overdetermined such that an intervention on a contributing actual cause is effective only if it is combined with a large number of additional interventions. Voting scenarios are an example for this kind of situation (Chockler and Halpern [2007]; Livengood [2013]). Suppose you are supporting a policy that is submitted for vote to a board constituted of eleven members. Suppose also that the policy will be put in place if there is a simple majority for it. You also know that currently there is a majority of eight members against your policy while it is supported by the remaining three board members. Each of the eight opposing board members would turn out as contributing actual causes of your policy not being put in place. Convincing any one of the board member will not make a difference to the prospects of your policy. You will have to convince at least three members.

Now, aren't such a functional argument and pluralism about actual causation separate issues? If one takes causal pluralism to mean extramural pluralism, then this may well be the case. After all, I spell out TAC, PAC, and CAC in one single theoretical framework: interven-

[8] As noted in Section 2, Halpern's 'modified' HP definition attempts to cover cases of symmetrical overdetermination by allowing conjunctions of variables as causes. But in light of the functional argument this is problematic because the definition does not distinguish whether one or several interventions have to be applied. For an extensive argument against conjunctions of variables as causes, see (Schaffer [2003]).

tionist causal models. However, intramural pluralism is supported by the functional argument. The functional approach evaluates whether and to which degree causal concepts facilitate intervention. I agree with extant accounts that intervention is facilitated by considerations about actual causes. But I argue that intervention is facilitated even better if we distinguish TAC, PAC, and CAC. Consequently, I argue, we should reject the assumption that there is a unified concept of actual causation and think in terms of TAC, PAC, and CAC instead.

Earlier I have stated that we tend to identify both the major and the sergeant as actual causes of the corporal's advancing. *Pace* Lewis's and Schaffer's intuition, this judgement is justified from the interventionist perspective. Both the major and the sergeant are potential targets for intervention if we want to prevent the corporal from advancing. It will not be sufficient to stop either one of the major and the sergeant to give their orders and to make them do nothing instead. We need to intervene on both to prevent the outcome. In this sense the trumping case is just like the situation of symmetrical overdetermination discussed in Section 2, where two contributing actual causes bring about an effect.

Unlike cases of symmetrical overdetermination, trumping also involves an asymmetry between the major and the sergeant, which is related to the fact that only the major is also a TAC. Again, this can be explained from the interventionist perspective. There is a straightforward sense in which the TAC is a better target for intervention: whereas an intervention on a mere CAC needs to be combined with other interventions, intervening on a TAC allows direct control over the outcome.

In order to see the benefit of distinguishing TAC and CAC more clearly let me contrast my account with Hitchcock and Knobe's ([2009]) functional account, which does not involve such a distinction. Hitchcock and Knobe's main concern are cases like the following. Suppose a fire would not have occurred if the refrigerator had not short-circuited or if not sufficient oxygen had been present. We typically identify the short circuit as actual cause and the oxygen as a mere background condition. What matters for such judgements, according to Hitchcock and Knobe, is difference-making in a 'normalized' version of the actual situation. In normal circumstances oxygen is present but no short circuits. In such circumstances oxygen doesn't make a difference with regard to there being a fire, but the short circuit does make a difference. Hitchcock and

Knobe also argue that such a selective concept of actual causation is useful. It enables us to identify those factors that are particularly suitable targets of intervention. Preventing short circuits, for example, is a better strategy for preventing fires than trying to avoid oxygen.

I agree that this is a viable approach to considerations of normality—considerations that are particularly important with regard to the problem of profligate causation discussed in the foregoing section. However, consider once again the trumping case. Suppose that both the major and the sergeant order the corporal to advance and shoot a captive civilian, a highly abnormal action (the relevant norm being a moral norm). Both the major and the sergeant are difference-makers in the 'normalized' version of the actual situation—the version where neither the major nor the sergeant order an immoral action but stay silent instead. So, both the major and the sergeant are equally suitable targets of intervention on Hitchcock and Knobe's account. Thus, unlike my account, Hitchcock and Knobe's account does not capture the fact that intervening on the major allows better control over the situation.

Similar considerations apply to the henchman case: both the gang leader and the henchman are difference-makers in the normalized version of the actual situation. But the gang leader is better suited as a target of intervention because manipulating the gang leader (a TAC) allows immediate control over the outcome, whereas an intervention on the henchman effects a change only if it is combined with an intervention on the gang leader. Finally, in the light bulb case my account identifies switch B as TAC, again indicating correctly that switch B allows better control over the outcome. For Hitchcock and Knobe's account suppose that the normal situation is one where the bight bulb is switched off. In any such situation the two switches are fixed in opposite positions ($A \neq B$). But then, again, either switch appears to be a difference-maker, meaning that both switches are equally suitable targets of intervention.

We shall now turn to path-changing actual causes (that are not also total actual causes). Again, we need (at least) two interventions in order to avoid the outcome. The important difference between path-changing actual causes (PACs) and other contributing actual causes that are not path-changing actual causes (CAC\PACs) is the relation between the interventions that have to be applied. In the case of CAC\PACs we need to apply two interventions that target independently active causal processes. In the case of PACs we need to combine a primary inter-

vention that targets the path-changing actual cause with a secondary intervention that eliminates a threat to the goal that did not exist (or at least was not as acute) before the primary intervention was applied. That is, the secondary intervention needs to be performed in order to counteract the adverse consequences of the primary intervention.

For example, in pre-emption cases, where the assassin in training is a path-changing actual cause of the victim's death, we need to apply two interventions. We need to apply a primary intervention on the assassin in training. But we also need to apply a secondary intervention on the supervising assassin because otherwise the supervisor would attack the victim. The difference to the overdetermination case is that the supervisor would attack the victim only as a result of the primary intervention.

This is an important difference for epistemic reasons. Elements of a set of CAC\PACs are more straightforwardly identified as threats to the desired outcome because each of them corresponds to an active causal process. Situations involving path-changing actual causes can be much more difficult to handle because sometimes we find out about the backup threats only through applying the primary intervention. This is because the backups are only activated as a result of the primary intervention.

This is even more problematic considering that sometimes intervening only on the path-changing actual cause can make the situation even worse. Suppose, for example, that the assassin in training is known to hurt her victims severely but does not kill them. Her supervisor, however, always hits her victims lethally. Suppose also that, given there is no interference, the trainee's attack pre-empts the supervisor's attack such that the victim is only hurt but not hit lethally. Intervening on the trainee's mission (but not on the supervisor) would then have detrimental consequences to our goal of saving the victim.

Safety mechanisms are another kind of example along these lines. Suppose you want to eat a piece of toast. However, as you put the bread into the toaster, the fuse is blown. You put the fuse back in but it keeps being blown whenever you start the toaster. You think you should intervene on the faulty fuse because you identify it as the actual cause of your not being able to enjoy your piece of toast. However, the fuse is only a path-changing actual cause. It pre-empts a short circuit and fire which would be an actual cause of your not enjoying the piece of toast,

and other consequences that may be much worse.

Moreover, consider complex policy decisions. Suppose the members of a city council want to make cycling safer and consider issuing a law that requires cyclists to wear a helmet.[9] Helmets are generally considered to increase cycling safety. However, wearing a helmet can also negatively affect safety because motorists are encouraged to leave less space when overtaking cyclists wearing a helmet (Walker [2007]). Suppose that before the bicycle law was issued Suzy used to cycle without a helmet. But now that the law is in place she wears a helmet while riding. A car overtakes while leaving little space, Suzy falls and due to her wearing a helmet she does not suffer head injuries. We are tempted to say that Suzy benefited from the helmet law because her wearing the helmet prevented severe head injuries. However, whether this is true depends on whether not wearing a helmet is a total actual cause or merely a path-changing (and not also a total) actual cause of an increased risk during cycling. If not wearing a bicycle helmet is a total actual cause, then Suzy seems to have benefited from the law. But if not wearing a helmet is only a path-changing actual cause then this is not necessarily so. It could be the case that without the helmet the motorist would have left more space and Suzy had not fallen in the first place.

7 Conclusion

In this article I have argued that we need to distinguish between three concepts of actual causation. The concept of total actual causation (TAC) amounts to straightforward counterfactual dependence. It is a widely used concept, but fails in situations that involve redundancy. In such situations we employ a concept of path-changing (PAC) or contributing actual causation (CAC). I have discussed three examples that have been troublesome for unified accounts of actual causation and I have shown that distinguishing between TAC, PAC, and CAC helps in these instances. I have also argued that the distinction is important from the perspective of an intervening agent. If $X = x$ is a TAC of $Y = y$ a simple intervention on X is sufficient to prevent $Y = y$. If $X = x$ is a PAC (but not a TAC), then a primary intervention on $X = x$ has to be combined with secondary interventions that counteract the adverse consequences of the

[9] This example is taken from (Hitchcock [2017]).

primary intervention. Finally, if $X = x$ is a CAC (but not a PAC), then an intervention on $X = x$ has to be combined with other independent interventions in order to prevent $Y = y$.

This being said, I do not claim that my account is exhaustive. Further research may show that other kinds of contexts require further causal concepts, beyond TAC, PAC, and CAC. A causal concept along the lines of Halpern and Pearl's concept of strong cause could be useful, for instance, in contexts where we want to sustain a certain outcome. In such circumstances we should be particularly interested in factors that guarantee the desired outcome independently of other factors—and this seems to be captured by the definition of strong cause.

Acknowledgements

Many thanks to Hasok Chang, Mathias Frisch, Christopher Hitchcock, Lucie White, and two anonymous reviewers for helpful comments. Funding for this research was provided by the Deutsche Forschungsgemeinschaft (254954344/GRK2073, 234743567/FOR2063).

Interdisziplinäres Zentrum für Wissenschafts- und Technikforschung
Bergische Universität Wuppertal
Wuppertal, Germany
efischer@uni-wuppertal.de

References

Cartwright, N. [1979]: 'Causal Laws and Effective Strategies', *Noûs*, **13**, pp. 419–37.

Cartwright, N. [2007]: *Hunting Causes and Using Them*, Cambridge: Cambridge University Press.

Chockler, H. and Halpern, J. [2004]: 'Responsibility and Blame: A Structural-Model Approach', *Journal of Artificial Intelligence Research*, **22**, pp. 93–115.

Danks, D. [2013]: 'Functions and Cognitive Bases for the Concept of Actual Causation', *Erkenntnis*, **78**, pp. 111–28.

Hall, N. [2004]: 'Two Concepts of Causation', in J. Collins, N. Hall and L. A. Paul (*eds*), *Causation and Counterfactuals*, Cambridge, MA: MIT Press, pp. 225–76.

Hall, N. [2007]: 'Structural Equations and Causation', *Philosophical Studies*, **132**, pp. 109–36.

Halpern, J. Y. [2016]: *Actual Causality*, Cambridge, MA: MIT Press.

Halpern, J. Y. and Hitchcock, C. [2010]: 'Actual Causation and the Art of Modeling', in R. Dechter, H. Geffner and J. Y. Halpern (*eds*), *Heurisitcs, Probability, and Causality: A Tribute to Judea Pearl*, London: College Publications, pp. 383–406.

Halpern, J. Y. and Hitchcock, C. [2015]: 'Graded Causation and Defaults', *British Journal for the Philosophy of Science*, **66**, pp. 413–57.

Halpern, J. Y. and Pearl, J. [2005]: 'Causes and Explanations: A Structural-Model Approach, Part I: Causes', *British Journal for the Philosophy of Science*, **56**, pp. 843–87.

Hitchcock, C. [2001]: 'The Intransitivity of Causation Revealed in Equations and Graphs', *The Journal of Philosophy*, **98**, pp. 273–99.

Hitchcock, C. [2007]: 'How to Be a Causal Pluralist', in P. K. Machamer (*ed.*), *Thinking about Causes: From Greek Philosophy to Modern Physics*, Pittsburgh, PA: University of Pittsburgh Press, pp. 200–21.

Hitchcock, C. [2017]: 'Actual Causation: What's the Use?', in H. Beebee, C. Hitchcock and H. Price (*eds*), *Making a Difference*, Oxford: Oxford University Press, pp. 116–31.

Hitchcock, C. and Knobe, J. [2009]: 'Cause and Norm', *The Journal of Philosophy*, **106**, pp. 587–612.

Knobe, J. and Fraser, B. [2008]: 'Causal Judgment and Moral Judgment: Two Experiments', in W. Sinnott-Armstrong (*ed.*), *Moral Psychology*, Vol. 2, Cambridge, MA: MIT Press, pp. 441–8.

Lewis, D. [1973]: 'Causation', *The Journal of Philosophy*, **70**, pp. 556–67.

Lewis, D. [2004]: 'Causation as Influence', in J. Collins, N. Hall and L. A. Paul (*eds*), *Causation and Counterfactuals*, Cambridge, MA: MIT Press, pp. 75–106.

Livengood, J. [2013]: 'Actual Causation and Simple Voting Scenarios', *Noûs*, **47**, pp. 316–45.

McDermott, M. [1995]: 'Redundant Causation', *British Journal for the Philosophy of Science*, **46**, pp. 523–44.

Menzies, P. [2004]: 'Difference-Making in Context', in J. Collins, N. Hall and L. A. Paul (*eds*), *Causation and Counterfactuals*, Cambridge, MA: MIT Press, pp. 139–80

Pearl, J. [2000]: *Causality: Models, Reasoning, and Inference*, Cambridge: Cambridge University Press.

Price, H. [2017]: 'Causation, Intervention, and Agency: Woodward on Menzies and Price', in H. Beebee, C. Hitchcock and H. Price (*eds*), *Making a Difference: Essays on the Philosophy of Causation*, Oxford: Oxford University Press, pp. 73–98.

Rosenberg, I. and Glymour, C. [2018]: 'Review of Joseph Halpern, *Actual Causality*', *BJPS Review of Books*, available at <www.thebsps.org/reviewofbooks/joseph-y-halpern-actual-causality/>.

Schaffer, J. [2000]: 'Trumping Preemption', *The Journal of Philosophy*, **97**, pp. 165–81.

Schaffer, J. [2003]: 'Overdetermining Causes', *Philosophical Studies*, **114**, pp. 23–45.

Schaffer, J. [2004]: 'Causes Need Not Be Physically Connected to Their Effects', in C. Hitchcock (*ed.*), *Contemporary Debates in Philosophy of Science*, Oxford: Basil Blackwell, pp. 197–216.

Spirtes, P., Glymour, C. and Scheines, R. [1993]: *Causation, Prediction, and Search*, New York: Springer.

Walker, I. [2007]: 'Drivers Overtaking Bicyclists: Objective Data on the Effects of Riding Position, Helmet Use, Vehicle Type, and Apparent Gender', *Accident Analysis and Prevention*, **39**, pp. 417–25.

Weslake, B. [unpublished]: 'A Partial Theory of Actual Causation'.

Woodward, J. [2003]: *Making Things Happen*, Oxford: Oxford University Press.

Woodward, J. [2014]: 'A Functional Account of Causation; or, a Defense of the Legitimacy of Causal Thinking by Reference to the Only Standard That Matters—Usefulness (as Opposed to Metaphysics or Agreement with Intuitive Judgment)', *Philosophy of Science*, **81**, pp. 691–713.