

Disjunctive Causes^{*}

I

There is an initial presumption against disjunctive causes. First of all, for some people causation is a relation between events. But, arguably, there are no disjunctive events, since events are particulars and thus they have spatiotemporal locations, while it is unclear what the spatiotemporal location of a disjunctive event could be.¹ More importantly, even if one believes that entities like facts can enter in causal relations, and even if there are disjunctive facts, it is still hard to see how disjunctive facts could be causes. Imagine, for instance, the following scenario. I have a gun filled with red paint and another gun filled with blue paint, and I fire both guns at my neighbor's white wall. A moment later, there is a graffiti on the wall and my neighbor notifies the police. He would have done so regardless of the graffiti's color, since all he cares about is the existence of a graffiti on his wall. Is it plausible to claim that a disjunctive fact is a cause of his notifying the police? In particular, is it plausible to claim that he notified the police because I fired the red-paint gun *or* the blue-paint gun (the thought being that my firing paint of either color *would have sufficed*)? It seems not. The police was notified because of the actual graffiti on the wall, and the actual graffiti on the wall is made of a certain pattern of colored patches. Imagine, that, as it turns out, there are patches of both colors on the wall. Then it seems that both my firing the red-paint gun and my firing the blue-paint gun were causes of my neighbor's notifying the police. In other words, my firing the red-paint gun and my firing the blue-paint gun jointly caused the outcome: each of them was a contributory cause of the outcome's occurrence. On the other hand, imagine that there are only patches of one color on the wall. Then it seems that my firing only one of the guns was a cause. Either way, the disjunction fails to be a cause: either my firing the red-paint gun was a cause, or my firing the blue-paint gun was a cause, or they were both causes, but their disjunction was not. This seems to generalize: for every outcome that could have been caused by more than one route, the relevant disjunctive fact (the fact that *some* route was

active) is not a cause of the outcome; instead, the outcome's causes are determined by the properties of the actual process issuing in the outcome's occurrence. Thus, it seems that causes cannot be disjunctive.

Now, in a few exceptional cases, some people are prepared to drop this presumption and to accept the existence of disjunctive causes. Two main types of case come to mind. First, overdetermination cases. Suppose that, improbably enough, two rocks hit a window at exactly the same time and the window shatters as a result. Suppose that each rock was independently sufficient for the shattering. What caused the shattering? If we were to say that both rocks did, we would be committed to a superabundance of causes (more causes than are needed to account for the fact that the window shattered). By contrast, if we were to say that the shattering was caused by the fact that *at least one* of the rocks struck the window—a disjunctive fact—the redundancy would be avoided. As a result, some people have thought that we should say that the disjunctive fact is a cause.² A second type of case in which someone might be prepared to accept the existence of disjunctive causes is causation by omission (or by absences in general). Omissions seem to be overly disjunctive in nature: they obtain just in case the world is one out of many different possible ways. For instance, my failure to go to the movies obtains just in case I stay at home or I go to a party or I go for a walk, etc. Thus, if omissions are causes, and if omissions are disjunctive, it follows that there are some disjunctive causes.³


Now, what we should say about cases involving overdeterminers and omissions is controversial. For one thing, some people find that there is nothing problematic about causal redundancy. If so, there is no motivation to posit the existence of disjunctive causes in cases of overdetermination. As for omissions, some people simply reject the possibility of their entering in causal relations. Still others accept the possibility of causation by omission but reject the claim that omissions are disjunctive (e.g., by identifying omissions with commissions). In any case the commitment to disjunctive causes is avoided.⁴

At any rate, setting aside exceptional and controversial cases like those involving overdeterminers and omissions, the presumption that causes are not disjunctive is quite general. In other words, the common thought is that, at least in ordinary circumstances, disjunctive facts cannot be causes. And a reason to believe this is that, as the graffiti case illustrates, it is hard to see how causes could be disjunctive in ordinary circumstances: it seems that, regardless of what the world is like, disjunctive facts cannot causally contribute to the occurrence of outcomes.⁵

In what follows I argue that this impression is misguided. I suggest that there is an intelligible way the world could be like for disjunctive facts to be causes in ordinary circumstances, and that the view that disjunctions are causes in those circumstances is supported by some intuitions about causation. I lay out the argument for disjunctive causes in the next section. Then, in section III, I explain how the argument can be generalized and, in IV, I examine some of its potential implications.

II

Let us start by imagining the following scenario. A runaway train is running on a track, where a person is trapped, up ahead. There is a switch and a side track. Someone flips the switch and as a result the train turns onto the side track—call this person *Flipper*. Unfortunately, the tracks reconverge after a while, where the person is trapped, and thus he still dies (in pretty much the same way, at around the same time). The situation is depicted by the following picture:



What were the causes of the person's death in this case? Clearly, we would say that they include: the train starting to run on the tracks, the driver's becoming unconscious, the person's foot getting stuck in the tracks, the train hitting the person, etc. On reflection, we would probably also identify other things as causes, in particular, events or conditions without which the death would not have occurred, but that are normally less salient to us, such as the tracks being in working order or the train's not falling apart as it travels down the track. All these, philosophers tend to think nowadays, are causes, although the contribution of the latter (which we could refer to as "background conditions") is sometimes unnoticed because it is too obvious to be worth mentioning.


Now, how about Flipper's action? Is his redirecting the train a cause of the person's death? On the face of it, it is not. It is clearly not one of the salient causes. Also, it does not seem to be one of the background conditions: it does not contribute something towards the outcome that is too unexciting or obvious to be worth mentioning. What the switching does is help determine the actual *route* to the outcome, and, intuitively, determining the route to an outcome is not tantamount to making a causal contribution to the occurrence of the outcome itself (however boring or obvious). To illustrate with another example: we feel that it would be wrong to say that an air stream that temporarily deflects a missile from its path (where the missile then resumes its original path) causes the target's destruction. This is so even if we think that the destruction of the target occurred due to many contributing causes (such as the missile's being launched, someone's giving the order to launch it, the presence of the target, the missile's never losing the necessary momentum to reach the target, etc.). For we do not think that the deflecting air stream was one of those contributory causes.

The judgment that Flipper's action is not among the causes of the death is quite generally recognized in the philosophical literature on causation and there have been several attempts to capture it.⁶ Recently, however, Ned Hall has argued that such judgment is misguided.⁷ Although this is not the place to offer a full defense of the intuition from possible objections (my main goal

here is to show that an argument for disjunctive causes can be built on that basis, *if* the intuition is to be trusted), I do think that we should aim to capture that intuition, and thus I will quickly review my main reasons for thinking this.

First, as I have already pointed out, the contribution of a switching event (like Flipper's redirecting the train) seems to be importantly different from even that of the least salient causes.⁸ Second, the intuition that switches are not causes seems to be importantly tied to the idea that causes are "difference-makers," which is a fundamental idea about causation—in particular, the reason Flipper appears not to be a cause is that it "fails to make a difference" in the relevant sense.⁹ And finally, the intuition that switches are not causes can do important theoretical work in grounding certain *moral* claims that we want to make about switching scenarios. For instance, we want to say that Flipper is not morally responsible for the person's death. Why is he not morally responsible? Intuitively, because he is not a cause of the death. Thus preserving the causal intuition allows us to ground attributions of moral responsibility on attributions of causal responsibility in the natural way.¹⁰

Now, an important feature of the train case, as I am conceiving it, is that the person would still have died if Flipper had failed to flip the switch (and maybe also that the redirection doesn't significantly alter the properties of the death). If the person would not have died had Flipper not flipped the switch, then, presumably, flipping the switch *would* be a cause of the death. For instance, imagine that the stretch of the main track where the two tracks come apart (call it S) is *disconnected*, as depicted by the following picture:



In this case, if Flipper had failed to flip the switch, the train would have derailed while passing through S and the person's life would have been spared. Thus flipping the switch *would* be a cause of the death. However, if the train would not have derailed if it had stayed on the main track because S is connected at the relevant time, then flipping the switch wouldn't be a cause of the death. Let this be claim 1:

C1: If S is connected and thus the train would have reached the person via the main track, then Flipper's action is not a cause of the person's death.¹¹

Now set Flipper aside for a moment. Imagine that S has been temporarily disconnected for repairs, and that someone reconnects it when the train passes through the switch—call him *Reconnector*. Is Reconnector's action a cause of the death of the person trapped on the tracks? Clearly, it depends: it depends on whether or not the switch was pulled. If the switch wasn't pulled and the train stayed on the main track, then the reconnection of S seems to be a cause of the death because, had it not been for the reconnection, the train would have derailed and the death would not have occurred. On the other hand, if the switch was pulled, then the train never had a chance to run on S. In that case, clearly, the reconnection of S is *not* a cause of the death (I am imagining that the pulling of the switch and the reconnection of S were causally independent events: neither occurred because the other occurred). Let this be claim 2:

C2: If the switch is (independently) pulled, and thus the train never runs on S, then Reconnector's action is not a cause of the person's death.

Again, this intuition is very strong: if the train never ran on S (and what the train actually did was not affected by S's reconnection in any way) then the reconnection had no effect on the person's fate.

On the basis of these claims, the argument for disjunctive causes is the following. Imagine a situation involving the train, the person tied to the track, and *both* Flipper and Reconnector. In particular, imagine that, when the train is about to pass through the switch, Flipper flips the switch and Reconnector independently reconnects S. The first premise of the argument is about Flipper. It says:

1. Flipper's action is not a cause of the death.

This follows from C1: given that Reconnector reconnected S, the train would also have reached the person if Flipper had not flipped the switch and the train had stayed on the main track. Thus, it follows from C1 that Flipper's action is not a cause of the death.¹²

The second premise is about Reconnector. It says:

2. Reconnector's action is not a cause of the death.

This follows from C2: given that Flipper flipped the switch, the train never ran on S. Thus, C2 entails that Reconnector's action is not a cause of the death.

Now, the third premise says:

3. Flipper and Reconnector *together* caused the person's death.

The justification for 3 is the following. If neither Flipper nor Reconnector had acted the way they did, then the train would have stayed on the main track, it would have derailed while passing

through S, and the person would have lived. Thus, the person died because of what the two of them did.

Now, what does it mean to say that Flipper and Reconnector *together* caused the death? It could mean one of two things. First, it could mean that the death happened partly because of what Flipper did *and* partly because of what Reconnector did. Second, it could mean that the death happened because *at least one of them* did what they did. The first suggestion is the “conjunctivist” suggestion, according to which the sum of the two actions, or a conjunctive fact, caused the death. The second suggestion is the “disjunctivist” suggestion, according to which a disjunctive fact caused the death. This is the fourth premise of the argument:

4. If Flipper and Reconnector together caused the death but their individual actions did not, then either the sum of the two actions (a conjunctive fact) caused the death or a disjunctive fact (the fact that either Flipper acted or Reconnector acted) caused the death.

Finally, the fifth premise says:

5. If either the sum or the disjunctive fact caused the death but the individual actions did not, then the disjunctive fact caused the death.

Note that, from 5 and the other premises, it follows that:

6. There are disjunctive causes.

Why believe in 5? I will mention two reasons. First, there is counterfactual dependence between the disjunctive fact and the death, but not between the conjunctive fact and the death.

There is counterfactual dependence between two actual events or facts C and E just in case, if C had not occurred/obtained, then E would not have occurred/obtained (or, according to the standard semantics for counterfactuals, if in the closest possible world(s) where C does not occur/obtain, E does not occur/obtain either). It is clear that the death counterfactually depends on the disjunctive fact: any nearby world where the disjunctive fact does not obtain is a world in which the person tied to the track lives. For it is a world where *both* Flipper and Reconnector fail to act in the way they do in the actual world, and, as a result, in any such world the train continues to run on the main track and derails while passing through S. By contrast, arguably, the death does not counterfactually depend on the sum of Flipper's and Reconnector's actions. There are nearby worlds in which the sum of the two actions does not obtain and in which the person tied to the track still dies. For instance, a world in which Flipper flips the switch but Reconnector does not reconnect S is a world in which the sum of the two actions does not obtain (given that one of those actions fails to obtain) but the person still dies (in that world, the train runs on the side track for a while, then on the main track again, and finally kills the person). Furthermore, that world seems to be closer to actuality than a world in which *neither* of the two actions obtains, which differs from the actual world in what both Flipper and Reconnector do. As a result, the death does not counterfactually depend on the sum of the two actions. So, if, as it is generally believed, the existence of counterfactual dependence is a strong indication of the existence of a causal relation, the disjunctive fact is a better candidate to be a cause than the sum of the two actions.¹³

A second reason to prefer the disjunctivist view to the conjunctivist view is that it is in general implausible to attribute causal powers to a sum if the sum did not get those causal powers from the parts. The relation between a sum and its parts is mereological composition: a sum is composed of its parts just like ordinary objects (say, rocks) are composed of atoms. Now, we would not say that a rock causes something, say, the shattering of a window, unless we believed that its atoms (at least some of them) causally contributed to the shattering. Similarly, it seems wrong to suggest that a sum can cause something without any of the parts causally contributing to

the same outcome. But this is what the conjunctivist view would have us say: according to the conjunctivist view, although Flipper's and Reconnector's individual actions are not causes, their sum is. This problem does not arise—or, at least, it doesn't arise with the same force—in the case of disjunctive facts. For the relation between a disjunctive fact and its disjuncts is not composition: a disjunctive fact is not made “out of” its disjuncts in the same way that a rock is made out of the rock-atoms.

To illustrate this point further, it is helpful to consider the case of omissions again. Intuitively, some omissions are causes. For instance, intuitively, my failure to save a child who is drowning in a pond a few feet away from me is one of the causes of the child's death. Now, suppose for a moment that omissions are disjunctive facts; thus, my failure to save the child is my either staying on the shore eating ice cream or my walking home or ... (so on and so forth). It is not implausible to suggest that the disjunctive fact (my omission) had causal powers that none of the disjuncts had. Intuitively, the child died because of what I failed to do, not because of what I did in its place. Suppose, for instance, that I was eating ice cream on the shore the whole time. It is plausible to suggest that it was not my eating ice cream *per se*, but only what it *entailed* (i.e. the fact that I failed to jump into the water to rescue the child) that caused the child's death. Hence, there is no general motivation for believing that, when (if) a disjunctive fact is a cause, at least one of its disjuncts must also be a cause.

This concludes my argument for disjunctive causes. As we have seen, the argument is based on a scenario where two individual events (Flipper's action and Reconnector's action) fail to be causes of a given outcome (as strongly suggested by intuitive judgments C1 and C2), but where the outcome happens as a result of the actions taking place together. In that case, the disjunctive fact is the best candidate for being a cause. In the next section I discuss how this argument can be generalized.

III

The basic fact that the case of Flipper and Reconnector helps to bring out is the following. Many times there is more than one possible route to an outcome. Now, *sometimes*, when different routes to an outcome are simultaneously viable, a fact about one of the routes makes a fact about the other routes causally irrelevant. In particular, if one route is actualized at the same time that an obstacle to another route is removed, then the actualization of the former can render the clearing up of the latter irrelevant, and the clearing up of the latter can render the redirection to the former irrelevant. This is what happens in the case of Flipper and Reconnector: the redirection to the side track makes the reconnection of the main track causally inefficacious, and, *vice versa*, the reconnection of the main track makes the redirection to the side track causally inefficacious. If either event had happened without the other, then that event would have been causally efficacious (if the reconnection had happened without the redirection, then the reconnection would have been a cause, and if the redirection had happened without the reconnection, then the redirection would have been a cause). But, when both events happen, they deprive each other of causal efficacy. Now, the argument continues, the outcome's occurring still depends on the existence of *some* viable route. Thus, the fact that some route was viable—a disjunctive fact—is causally relevant to the outcome. Hence, there are disjunctive causes.

To clarify: I do not mean to suggest that the argument for disjunctive causes generalizes to *every* case where there is more than one potential route to an outcome. This is clearly false. So-called “preemption” cases illustrate this: in a case of preemption, there are two possible routes to an outcome but one of those routes *preempts* the other, this is to say, it renders the other route causally inefficacious by exerting its own causal efficacy. For example, if two bullets are fired at a target and one of the bullets hits the target first, the first bullet preempts the second bullet. In this case, the outcome fails to have a disjunctive cause; rather, the first bullet is a cause and the second bullet is not.

Under what conditions, then, does the existence of more than one potential route to an outcome give rise to a disjunctive cause? The answer is: only in cases where there is an event that

accounts for the switch in routes without, intuitively, being a cause of the outcome. The train scenario is a case of this sort: the two tracks are alternative routes to the outcome (the person's death) and, intuitively, all that the pulling of the switch does is to redirect an ongoing threat from one path to the other, without thereby being a cause of the outcome. In other words, the pulling of the switch determines the route to the death without being a cause of the death. In contrast with the train case, in a case of preemption, where, e.g., I shoot the first bullet at the target and someone else shoots the second bullet that then soars through empty space, intuitively I do cause the outcome (I do not merely redirect a preexisting threat).

The contrast, then, is between "switching" cases and "preemption" cases. However this difference is to be accounted for (I will not attempt to do this here)¹⁴, the suggestion is that the argument for disjunctive causes only generalizes to cases with the switching structure. The general form of the argument is as follows. Suppose that there are two (or more, but let us stick with two) potential routes to an outcome (say, A and B). Then, *in cases with the switching structure*, changing from A to B is not causally contributing to the outcome. Also, A's being viable does not causally contribute to the outcome if A is not used. However, the outcome would not have occurred if the switch from A to B had not taken place and if A had not been viable. Hence, since the disjunctive fact is a better candidate cause than the sum of the two events, it (the fact that either there was a switch from A to B or A was viable) is a cause of the outcome's occurring. Therefore, the outcome has a disjunctive cause.

Now we can see where the initial presumption against disjunctive causes went wrong. As I have pointed out, that presumption is based on the impression that there is no way the world could be like for causes to be disjunctive (at least in ordinary circumstances). I illustrated this impression with the graffiti example: even if the police would have been notified had the graffiti been made out of blue or red paint (or both), there is no possible state of affairs in virtue of which the cause of the police being notified could be a disjunctive fact (e.g., the fact that I fired a red-paint gun or a blue-paint gun at the wall). This is so because, if we look at the actual process

issuing in the creation of the graffiti and thus in the police being notified (that is to say, if we follow the actual traces of blue and red paint), we will be able to determine that the cause was the firing of one gun (if there are only traces of one color) or the firing of both (if there are traces of both colors), but the fact that *either* gun was fired would never fill the bill.

However, this reasoning fails to generalize to cases with the switching structure. An important difference between the graffiti case and switching cases is the following. Imagine that, as it turns out, the graffiti is only made out of blue paint. In that case, the mere fact that the graffiti *could* also have been made out of red paint (in which case the police would still have been notified) is not sufficient to conclude that my firing the blue-paint gun was not a cause of the police being notified. In contrast, in the train case (and in switching scenarios in general), the mere fact that there was an alternative route is sufficient to rob the event of the redirection of its causal powers. Intuitively, the redirection is not a cause given that there was an alternative route, even if that route is never actualized. That is the distinguishing mark of a switching case: in a switching case, the status of a merely potential route can affect the causal status of the redirection to the actual route. As a result, there *can* be disjunctive causes. We just could not see it because we were focusing on cases with the wrong causal structure.

Now, someone might worry that the Flipper and Reconnector case has taught us nothing new. In particular, someone might argue that it is just another case of overdetermination, and that the argument from overdetermination cases to disjunctive causes is—as I have acknowledged in section I—one that has already been recognized. Why would one be tempted to think that the Flipper and Reconnector case is a case of overdetermination? Because the switching and the reconnecting behave at least *like* overdeterminers in that, if each had taken place without the other, the outcome would still have occurred. If the reconnection had happened without the switching, or if the switching had happened without the reconnection, the death would still have occurred. Thus, one might think that the death was overdetermined by the two events.

However, the switching and the reconnecting differ from typical overdeterminers in at least two main ways. First, in a paradigm case of overdetermination (such as two rocks simultaneously hitting a window and making it shatter, or two bullets simultaneously entering a person's heart and causing that person's death), there are two easily distinguishable processes leading up to the outcome, and both of those processes are equally "active." In the train case, by contrast, we would not ordinarily say that there are two active processes, but only one: the actual trajectory of the train (if there had been two trains instead of one, each train running on one of the tracks, and if the two trains had reached the person at the same time, then that would have been a standard case of overdetermination). So, at the very least, the switching and the reconnecting are not *typical* overdeterminers. Whether in the end we should count them as overdeterminers is a separate matter, which I prefer to set aside here. The bottom line is that, even if they are overdeterminers, they are a special variety of overdeterminers, and as such they can supply a new set of reasons to believe in disjunctive causes.

Another important way in which the switching and the reconnecting differ from typical overdeterminers is that they need not occur simultaneously for them to give rise to the argument for disjunctive causes. Typical overdetermination cases involve some sort of unlikely coincidence at some point along the chain leading to the outcome (for instance, two rocks striking a window at exactly the same time, or two bullets entering a person's heart at exactly the same time). But this need not be the case with the switching and the reconnecting for the argument for disjunctive causes to work. To be sure, there are timing restrictions. In particular, one might argue that S had to be connected by the time the switch is actually pulled. Or one might argue that S had to be connected by the time the train *would* have reached it if the switch had not been pulled. Now, regardless of how exactly we should understand this restriction, it certainly does not require anything like a striking coincidence. In particular, it would be too strong to require that the reconnection and the switching happen at the same time. Suppose, for instance, that the reconnection happens a few seconds *before* the switching. In that case the argument for

disjunctive causes would still have the same force: the reconnection would make the switching causally inefficacious and, in turn, the switching would make the reconnection causally inefficacious.^{15, 16}

In sum, the Flipper and Reconnector case differs from typical overdetermination cases in that it does not involve two causal paths to an outcome both of which go to completion, and which go to completion at exactly the same time. So, again, even if it is a case of overdetermination, it is an atypical example, and as such there are interesting lessons to be learned from it.¹⁷

IV

Suppose that we are persuaded by the argument for disjunctive causes. What consequences does this have? In what follows I discuss three potential implications.

First, an obvious implication concerns the debate about the “causal relata.” What kinds of things can enter in causal relations? If we allow for disjunctive causes, then, unless we are prepared to accept the existence of disjunctive events, we should be open to the possibility that other things besides events can enter in causal relations—probably facts. (This might not be a surprise, if we already believed that, say, omissions can enter in causal relations.)

A second, less obvious consequence concerns the concept of moral responsibility and its relation to causation. Imagine a scenario where we would want to blame Flipper and/or Reconnector for the death of the person on the tracks. Imagine, for instance, that both Flipper and Reconnector had reasons to refrain from doing what they did (for instance, because they didn't know what the other would do, and thus they should have tried to prevent the death from happening by not doing their part). In that case the argument for disjunctive causes would have us say that, although each of them bears some moral responsibility for the outcome, neither causes it (more precisely, neither causes the outcome individually, although a disjunctive fact involving them does). If so, we should probably revisit some common beliefs about the relationship

between moral responsibility and causation. For, according to the received view of their relationship, being morally responsible for an outcome requires, at the very least, causing it. In contrast, the argument for disjunctive causes suggests that it is possible to be responsible for something without being a cause in the standard sense.¹⁸

Finally, it is worth examining the bearing that the argument for disjunctive causes has, if any, for an argument for the autonomy of the mental from the physical. Roughly, the argument for the autonomy of the mental goes as follows.¹⁹ Presumably, the mental is “multiply realizable” by the physical (this is to say, at least some mental events have different possible physical “realizers,” or supervenience bases). But then this means that, for at least some mental events, there are no single physical events that could be identified with them across the board. So, unless we are prepared to identify mental events with disjunctive events (the disjunctions of all the different possible realizers of those mental events), we should probably conclude that mental events are not identical to physical events. Moreover, the argument would continue, the option of identifying mental events with disjunctions of physical realizers is problematic. Here the reasons offered might diverge or multiply. One reason could be that there are no disjunctive events. A related reason could be that there might be an infinite number of physical realizers for a given mental event, but that there are no infinite disjunctions, or events corresponding to infinite disjunctions. Another type of reason could be that, even if there were disjunctive events, or disjunctive events of the relevant type, they would not be able to enter in causal relations since there are no disjunctive causes (or that, even if there were disjunctive events, they could not occur in laws since generalizations about disjunctions are not sufficiently law-like). Note that this argument is likely to extend to subvenient/supervenient events in general; the case of the physical and the mental is simply an instance of that more general argument, but it is helpful to focus on just one instance.

The question that arises, then, is: Can the argument for disjunctive causes be used to counter the argument for the independence of the mental from the physical? In particular, could it

offer reasons to identify mental events with disjunctions of physical realizers on the grounds that those disjunctions could be causally efficacious, after all? In this case, my answer is negative: the type of scenario on which the argument for disjunctive causes was based is sufficiently different from the case of mental/physical events to not warrant drawing such an implication. In particular, given those differences, even if the argument for disjunctive causes succeeds in showing that there are *some* disjunctive causes, it has no tendency to show that disjunctions of *physical realizers* are causes. Thus it does not serve the purpose of defending the identification of mental events with disjunctions of physical realizers.

Let me explain. A main difference between the case of Flipper and Reconnector and the case of a mental event and its physical realizers is the following. In the case of Flipper and Reconnector, as I have pointed out, it is crucial that both actions were actual. If Flipper had acted without Reconnector doing his part, or if Reconnector had acted without Flipper doing his part, the outcome would not have been disjunctively caused. This is so because a main feature of the Flipper and Reconnector case, and one *in virtue of which* I argued that we should say that the outcome was disjunctively caused, is that the two actions are “mutually neutralizing” (i.e. they deprive each other of causal efficacy). But, for them to neutralize each other, they must both obtain. In other words, the argument for disjunctive causes that I have presented here crucially relies on a type of scenario where both (all) of the disjuncts are present. For each disjunct is needed to neutralize the causal powers of the other disjuncts (in particular, if only one of the disjuncts were present, then that disjunct would be a cause and the disjunction itself would not). As a result, the argument offers no reason to believe that causes can be disjunctive when only one disjunct is present.

Now, it is likely that, out of the potential physical realizers of a mental event, at most one of them will be actual at any given time. At least this is the way in which people normally think of the different possible physical realizers: as mutually exclusive physical states, or as states that are likely to be realized by organisms of different types (say, humans and Martians). This makes

the case of mental events and their physical realizers importantly different from the Flipper and Reconnector case. In virtue of this fact, the argument for disjunctive causes is not likely to have significant implications for the case of the mental and the physical.²⁰

V

I have argued that the initial presumption against disjunctive causes is misguided, and that there is a persuasive argument for disjunctive causes. I have also discussed some potential implications of that argument. Whereas I noted that it has important implications for the debate over the causal relata and for the relationship between causation and moral responsibility, we should not expect it to have significant consequences for an argument that the mental is autonomous from the physical.

Carolina Sartorio

University of Wisconsin-Madison

* For valuable discussion and comments, thanks to Ellery Eells, Martha Gibson, Daniel Hausman, Russ Shafer-Landau, Lawrence Shapiro, Alan Sidelle, Stephen Yablo, the audience at the 2005 Pacific APA (in particular, my commentator, Brad Armendt), the members of the metaphysics & epistemology reading group at the University of Wisconsin-Madison, two editors of this Journal, and especially Juan Comesaña.

¹ For an argument of this type against disjunctive events, see David Lewis, "Events," in *Philosophical Papers, Volume II* (New York: Oxford, 1986), pp. 241-69.

² See, e.g., Jonathan Bennett, *Events and their Names* (Indianapolis: Hackett, 1988); John Mackie, "Causes and Conditions," in E. Sosa and M. Tooley, eds. *Causation* (New York: Oxford, 1993), pp. 33-55; D.H. Mellor, *The Facts of Causation* (London: Routledge, 1995).

³ David Lewis discusses this view on omissions in Postscript D to "Causation," in *Philosophical Papers, Volume II* (New York: Oxford, 1986), pp. 189-93. Other cases where one might be led to think that

disjunctions are causes are “preemptive prevention” cases (these cases also involve absences, although as intermediaries in the chain; see n. 17 below).

⁴ For an argument that overdetermination is widespread, see Jonathan Schaffer, “Overdetermining Causes,” *Philosophical Studies* 114 (2003): 23-45. For an argument that omissions cannot enter in causal relations, see Helen Beebe, “Causing and Nothingness,” in J. Collins, N. Hall, and L.A. Paul, eds. *Causation and Counterfactuals* (Cambridge: MIT, 2004), pp. 291-308.

⁵ Alan Penczek appears to be an exception (see his “Disjunctive Properties and Causal Efficacy,” *Philosophical Studies* 86 (1997): 203-19). Penczek seems to believe that any ordinary case where an outcome could be caused in one of two ways might be a case where the outcome has a disjunctive cause. But I find his argument unconvincing. Penczek would argue, for instance, that if I shoot a bullet that kills a person but there was a backup assassin waiting in reserve, then my shooting is not a cause because the relevant counterfactual dependencies fail to obtain (in particular, the death would still have occurred if I had not shot). Thus, since we do not want to say that the outcome is uncaused, we should say that a disjunctive condition (presumably, the fact that at least one of us shot) is the cause. However, if anything, cases of this type are a reason to believe that counterfactual accounts of causation are wrong, not a reason to believe in disjunctive causes.

⁶ See William Rowe, “Causing and Being Responsible for What is Inevitable,” *American Philosophical Quarterly* 26, 2 (1989): 153-59; Laurie Paul, “Aspect Causation,” *Journal of Philosophy*, 97, 4 (2000): 235-56; Stephen Yablo, “De Facto Dependence,” *Journal of Philosophy* 99 (2002): 130-48; Carolina Sartorio, “Causes as Difference-Makers,” *Philosophical Studies* 123, 1-2 (2005): 71-96.

⁷ Ned Hall, “Causation and the Price of Transitivity,” *Journal of Philosophy* 97, 4 (2000): 198-221 (Hall calls cases of this type “interactive switches,” given the existence of a physical connection between the switching event and the outcome). Hall seems to now have changed his view on these matters (personal communication).

⁸ In his *op. cit.*, Hall argued that switching events are boring causes on a par with background conditions. His argument went as follows. The switching’s contribution is like the contribution of the existence of the side track a day earlier (a boring cause). For the relation between the switching and the setting of the switch as the train passes over it is just like the relation between the existence of the side track the day before and

its existence today, when the train travels down it; in addition, the relation between the setting of the switch as the train passes over it and the outcome is just like the relation between the existence of the side track today and the outcome. Hall seems to be thinking in terms of physical connectedness: just like we can trace a continuous physical process linking the side track's presence the day before to the outcome today (via the presence of the track when the train passes over it), we can trace a similar process linking the switching event to the outcome (via the setting of the switch when the train passes over it). However, unless we assume that physical connectedness is sufficient for causation (which is debatable), this fails to show that the two relationships are the same in all relevant respects. In particular, as I have pointed out, a main feature of the switching event is that what it does can be seen as the deflection of an ongoing process (in virtue of which the switching helps to determine the actual route to the outcome, but this is all it does). In contrast, the presence of the side track the day before plays no similar deflecting role.

⁹ In what sense, exactly, does the switching fail to make a difference? One might think that it fails to make a difference in that the outcome would still have occurred in its absence; however, the existence of “preemption” cases shows that it is possible for an event to be a cause of an outcome even if the outcome would still have occurred in the absence of the event. In my *op. cit.*, I suggest that causes are difference-makers in this other sense: if an event is a cause, then its absence would not have been a cause of the same outcome. The switching fails to make a difference in this sense, for, if we were to count it as a cause, then, had it not occurred, we would have had to count its absence as a cause too.

¹⁰ A clarification is in order: as we will see later (see section IV below, and, in particular, n.18), I think that the fact that Flipper failed to cause the death can only be *part* of the explanation of his lack of moral responsibility; still, I think that it is an important part of the explanation. By the way, note that, although Flipper's lack of moral responsibility seems importantly connected with his alleged lack of causal responsibility, our intuition that he is not a cause is not “tainted” by the moral judgment, but it is a purely metaphysical intuition. For we still think that the redirection fails to cause the death if what redirects the train is, say, a gust of wind instead of a moral agent.

¹¹ Notice that I am not claiming that it is *generally* true that, when a person's death would still have occurred if one had acted differently, then one's action does not cause the death. The fact that the death

would still have occurred if Flipper hadn't acted clearly is part of the explanation of why he fails to be a cause, but it is not the whole explanation (given the existence of preemption cases; more on this later).

¹² One might think that the exact time of the reconnection matters to whether Flipper is a cause. I consider this possibility in section III.

¹³ Note that the argument does not rest on the claim that counterfactual dependence *is* causation (which, as has often been pointed out, is most likely false) but only on the weaker claim that there is a strong *correlation* between counterfactual dependence and causation.

For an alternative view of the relevant notion of similarity among possible worlds, see David Lewis, Postscript E to "Causation," in *Philosophical Papers, Volume II* (New York: Oxford, 1986), pp. 211-2. There Lewis claims that the closest world to one where a sum A&B does not obtain is one where neither A nor B obtain, or, in other words, one where the event has been completely "excised" from the world, which includes removing all of its parts. But it is unclear that this is how we ordinarily evaluate counterfactuals. Recall that, in the case of Flipper and Reconnector, the two actions are independent. Hence, imagining that one is absent should not lead us to imagine that the other one is absent too. At any rate, the point would remain that, whereas the outcome *clearly* depends on the disjunctive fact, this is not true of the conjunctive fact.

¹⁴ This is the type of causal difference that the authors mentioned in n.6 aimed to capture.

¹⁵ Recall that we are assuming that the switching and the reconnecting are causally independent events. How exactly to parse the timing restriction is an interesting problem. What if the reconnection happens *after* the switching? In his *op. cit.*, Yablo develops a theory according to which, very roughly, meeting a "need" is sufficient for being a cause. Yablo's theory seems to entail that the switching *is* a cause if the reconnection happens after the switching (although not otherwise) because then the switching meets a need for the death. But I am not sure that this is the right result. Suppose that there is a flawless mechanism already on its way to reconnect S, and it will do so before the train reaches S. It is somewhat implausible to suggest that the switching is only a cause if the mechanism reconnects S before the switching takes place. Clearly, whether the mechanism would operate in time for the train not to derail matters, but whether it would do so before or after the switch is flipped seems otherwise irrelevant.

¹⁶ Given that the case of Flipper and Reconnector does not require an improbable coincidence, it would still be an “ordinary” case even if it were a case of overdetermination (in the sense I have used the word “ordinary” in section I). There might be other “atypical” overdetermination cases that do not involve improbable coincidences. For instance, election cases where a candidate wins by more than one vote: arguably, in those cases the outcome of the election is overdetermined but it does not involve any unlikely coincidence. Also, if a 250-pound man steps on a scale that only registers up to 200 pounds, one might argue that what the scale registers is overdetermined, but it does not involve any coincidence. (In his *op. cit.*, Schaffer argues that overdetermination without improbable coincidences is very common and that it happens whenever, e.g., a rock hits a window with more momentum than was necessary to make it break.) Now, as far as I can see, none of these other cases support an argument for disjunctive causes—at least not an argument of the type I have put forth here.

¹⁷ Someone might think that the Flipper and Reconnector case resembles a “preemptive prevention” case more than an overdetermination case (for discussion of preemptive prevention cases, see Michael McDermott, “Redundant Causation,” *British Journal for the Philosophy of Science* 46 (1995): 523-44, and John Collins, “Preemptive Prevention,” *Journal of Philosophy* 97, 4 (2000): 223-34.). Suppose that a catcher catches a ball that was flying towards a window; however, a second catcher was standing between the first catcher and the window and would have caught the ball otherwise. What caused the window to remain intact? Clearly, the second catcher did not do anything. But, given the presence of the second catcher, the window was in no danger of breaking, so maybe the first catcher was not a cause either (this intuition is reinforced if we substitute a solid wall for the second catcher). So maybe the cause was a disjunctive fact involving the two catchers? (The thought being that, if neither catcher had been present, the window would have shattered.) There are certainly similarities between preemptive prevention cases and the Flipper and Reconnector case. But there are also important differences. One main difference is that preemptive prevention cases involve absences along the causal chain—e.g., the ball’s not hitting the wall—whereas the Flipper and Reconnector case does not. This makes the Flipper and Reconnector case an “ordinary” case (as defined in section I).

¹⁸ Typically we think that an agent causally contributes to an outcome just in case one of the agent’s individual actions or omissions causally contributes to the outcome. In “How to be Responsible for

Something without Causing It” (*Philosophical Perspectives* 18 (2004): 315-36) I lay out the received view about the relationship between moral responsibility and causation, and I argue against it on the basis of an independent example. I then argue that moral claims are still dependent on causal claims, but not in the way specified by the received view.

¹⁹ See, e.g., Jerry Fodor, “Special Sciences,” *Synthese* 28 (1974): 97-115.

²⁰ In this sense, the case of the physical and the mental is more similar to the case of omissions. For, if omissions were disjunctive entities, they would be disjunctions of mutually exclusive disjuncts (e.g., my failure to do A would be the disjunction of all the things that I could have done instead of doing A).