

Actuality and Responsibility

CAROLINA SARTORIO

University of Arizona

sartorio@arizona.edu

Actual-sequence views of responsibility are views according to which moral responsibility is a function of actual sequences, histories, or ancestries. In recent years these views have acquired much popularity as an attractive kind of compatibilist answer to the problem of determinism and the freedom of the will. But what does it mean to say that responsibility is ‘a function of the actual sequence’? In this paper I examine different possible ways to cash out this idea. I show that one of them is immune to important objections to which the others are prey. This motivates a type of actual-sequence view that is unorthodox in two main respects. First, it understands the expression ‘the actual sequence’ in a way that is different from the way in which it seems typically to have been understood. Second, on this view, non-actualized possibilities of a certain kind are always relevant to the moral responsibility of agents.

Following Frankfurt’s lead (Frankfurt 1969), many philosophers seem to believe that moral responsibility is a function of actual sequences, histories, or ancestries (a main example is Fischer; see, for instance, Fischer 2006b). According to this view, an agent’s responsibility for, say, an act is fully determined by the way in which the agent actually came to perform the act. What could have happened but did not happen—in particular, whether the agent could have acted in any other way (whether the agent had ‘alternative possibilities’ of action)—is irrelevant to the agent’s responsibility.¹

I will use the expression ‘actual-sequence views of responsibility’ to refer to views according to which an agent’s responsibility for something is a function of the actual sequence issuing in it. In recent years, actual-sequence views of responsibility have been associated with a particular solution to the problem of determinism and the freedom of the will (sometimes referred to as ‘semi-compatibilism’) according to which responsibility is compatible with the absence of alternative

¹ In his introduction to his anthology on free will, Watson writes: ‘Frankfurt’s essay teaches us that what matters for responsibility is the actual explanation of our behaviour’ (Watson 2007, p. 7).

possibilities for action, and so responsibility is not incompatible with determinism even if determinism rules out alternative possibilities for action. On these views, responsibility is compatible with the absence of alternative possibilities because an agent is responsible to the extent that the actual sequence leading to his act has certain properties; what the agent could have done but did not do is irrelevant to the agent's responsibility, since it is a mere possibility. Thus, a presumed advantage of these views is that they allow us to preserve the idea that people can be responsible under the assumption of determinism in the face of incompatibilist arguments that seemingly show that determinism rules out the ability to do otherwise (such as van Inwagen's famous 'consequence argument'; van Inwagen 1983).

What, though, does it mean to say that responsibility is 'a function of the actual sequence'? This idea has not been given a clear and precise interpretation. In this paper I will examine different possible ways to cash it out. At first sight, all of these possible interpretations seem to run into serious problems. However, I will argue that one of them can survive the main objections, when it is interpreted in the right way. This will motivate a kind of actual-sequence view that is unorthodox in two main respects. First, it understands the expression 'the actual sequence' in a way that is different from the way in which it seems typically to have been understood. Second, on this view, mere possibilities are relevant to an agent's responsibility, even if the agent's responsibility is a function of the actual sequence. In fact, on this view unactualized possibilities of a certain kind are *always* relevant to an agent's responsibility.

1.

What could be meant by the claim that responsibility is a function of the actual sequence? A quite natural interpretation of that claim is this:

- (1) *Irrelevance of causally inefficacious factors:*
When an agent is responsible for X, only factors that are links in the actual causal sequence issuing in X ground the agent's responsibility for X

Note that a specific consequence of (1) is this:

- (2) *Irrelevance of mere possibilities:*
When an agent is responsible for X, unactualized possibilities do not ground the agent's responsibility for X

For any (actual) X, unactualized possibilities, or ‘mere’ possibilities, are never links in the actual causal history of X. Hence, according to (1), they cannot ground the agent’s responsibility for X. In particular, possible scenarios where an agent acts differently from the actual way in which he acts are irrelevant to the agent’s responsibility for his act in the actual world. Thus, whether the agent had alternative possibilities of action is irrelevant to the agent’s responsibility for his act.²

Another way to try to characterize the claim that responsibility is a function of the actual sequence is in terms of a supervenience claim:

(3) *Supervenience:*

There cannot be a difference in responsibility for X without a difference in the actual causal sequence issuing in X

For example, there cannot be two scenarios involving agents who perform the same act as a result of the same causal sequence of events and where one of the agents is responsible for his act but the other is not. In other words: if the agents came to perform the same act in exactly the same way, one is responsible if and only if the other one is.³

Frankfurt endorsed theses (1) through (3) (at least when X is an act or an omission; more on this below).⁴ He did this as part of his attack

² It is common to think that the only possibilities that could matter for responsibility are possibilities that are accessible to the agent. Given that (2) rules out the relevance of all possibilities, it rules out the relevance of accessible possibilities as well as that of inaccessible possibilities.

³ A supervenience claim is only a covariation claim. Hence, in order to capture the idea that responsibility is *grounded* in actual sequences, perhaps we should strengthen the supervenience claim into a ‘superdupervenience’ claim (basically, a claim to the effect that the differences of the subvenient kind *ground* the differences of the supervenient kind; see Schiffer 1987). The superdupervenience claim would read:

(3*) *Superdupervenience*

There cannot be a difference in responsibility for X without a difference in the actual causal sequence issuing in X; moreover, when there is a difference in responsibility, the corresponding difference in the actual sequence grounds the difference in responsibility.

Here I will focus on the supervenience claim, but everything I say about that claim applies, *mutatis mutandis*, to the superdupervenience claim.

⁴ In Frankfurt 2006, Frankfurt suggests that reference to the actual *causal* sequence might be too narrow; one should use ‘actual *explanatory* sequence’ instead. This is because he thinks that there are factors that explain why an agent performed an act although they do not cause the act. I do not see the motivation to say this (see, in particular, n. 10 below). Hence I will proceed as if the actual explanatory sequence were just the causal sequence.

on the ‘principle of alternative possibilities’, that is, the principle according to which responsibility requires alternative possibilities of action (see Frankfurt 1969; Frankfurt used the expression ‘alternate possibilities’ instead of ‘alternative possibilities’). Frankfurt reasoned as follows. Start with a relatively uncontroversial case of responsibility where an agent is responsible for an act that he could have avoided and that he decided to perform completely ‘on his own’ (say, as a result of having freely considered reasons for and against performing it and after having weighed those reasons against each other). Frankfurt claimed that we hold the agent responsible in this case because he was moved to act by his own reasons. Now add a sneaky neuroscientist waiting in the wings, ready to intervene by making the agent decide to act in the relevant way if he does not do this on his own. Frankfurt argued that, if the agent still acts moved by his own reasons, then he is equally responsible in this case (scenarios of this kind have since then been called ‘Frankfurt-style’ scenarios). For if what made the agent responsible in the ordinary (neuroscientist-free) case was the fact that he acted for his own reasons, then, given that he also acts for his own reasons in the Frankfurt-style scenario, he is equally responsible in that case. Any difference in responsibility between a Frankfurt-style scenario and an ordinary scenario would have to be due to a difference in the actual sequence leading to the act (thesis (3)). Factors that do not play a role in causing the agent to act (such as the neuroscientist) are irrelevant to the agent’s responsibility for his act (thesis (1)). In particular, unactualized possibilities — such as what would have happened if the agent had not decided to act on his own (the fact that the neuro-scientist would have intervened in those circumstances) — are irrelevant to the agent’s responsibility for his act (thesis (2)).⁵

Frankfurt argued that similar reasoning applies to responsibility for omissions (Frankfurt 1994). And, although Frankfurt did not himself discuss the case of outcomes, it is easy to extend the same style of reasoning to responsibility for outcomes, in the following way. Imagine that, as a result of the agent’s deciding on his own to

⁵ Frankfurt suggested that a different principle of alternative possibilities might be true, one according to which an agent is not responsible for an act if he acted only because he could not have done otherwise (only as a result of those factors in virtue of which he could not have done otherwise). For example, if an agent was coerced into doing something, and he did it only because of the threat that was launched against him, he is not responsible (Frankfurt 1969). Note that, according to this principle, mere possibilities are never relevant to an agent’s responsibility. Factors in virtue of which the agent could not have done otherwise are irrelevant to the agent’s responsibility unless they are actual causes of the agent’s act.

punch someone in the face, his victim is now in pain. Then, one could argue, the agent is responsible for his victim's pain even if his victim could not have failed to be in pain (say, because the agent could not have failed to punch him, in light of the neuroscientist's presence). The neuroscientist's presence is irrelevant to the agent's responsibility for the outcome because it was not a link in the actual sequence leading to that outcome. Mere possibilities are irrelevant, since they are not links in the actual sequence. For similar reasons, responsibility for an outcome supervenes on the actual sequence issuing in the outcome.

A problem arises at this point, however. It seems that two agents could bring about an outcome in exactly the same way (as a result of the same causal sequence) but one be responsible for the outcome — because he could foresee (or should have foreseen) that the outcome would follow — and the other one not be responsible — because he could not have foreseen (or it is not the case that he should have foreseen) that the outcome would follow.⁶ For example, imagine that I press a button because I know that it will make a squeaky sound when I press it, and I find the squeaky sound amusing. As it turns out, pressing the button also results in an explosion that causes some harm. I was not aware of (and it is not the case that I should have been aware of) the fact that pressing the button was likely to have such bad consequences. Imagine that you, on the other hand, were aware of this fact. Still, you pressed the button, just like me, because you wanted to hear the squeaky sound. Unlike me, you are responsible for the harm, and this is partly in virtue of the fact that you could foresee that the harm would follow. But your belief that the harm would follow was not a link in the causal chain leading to your act (it did not cause you to press the button). This suggests that an agent's responsibility for an outcome is not just a function of the actual sequence of events leading to the outcome: other factors can also ground the agent's responsibility.

I do not think that this is an insurmountable problem, though. It is common to distinguish the so-called metaphysical conditions for responsibility from the so-called epistemic conditions. Typically, the metaphysical conditions are regarded as the conditions that have to obtain for an agent to have the relevant kind of control, or freedom, necessary for responsibility; the epistemic conditions, on the other

⁶ Perhaps this is true of acts and omissions too, to the extent that agents may not be able to foresee that they will perform a certain act or omission.

hand, are of a quite different nature: they concern the existence of certain mental states such as the agent's beliefs, intentions, etc.⁷ Ginet (2000), for example, argues that what has to be added to the metaphysical conditions (in the case of blameworthiness for a harm) is a condition stating, very roughly, that the agent knew or should have known that the harm would or might occur as a result of his act. This would explain why you are responsible for the harm by pressing the button and I am not.

Now, I do not mean to suggest that what categories of things belong in the set of the epistemic conditions for responsibility is totally clear.⁸ For example, it is not fully clear what the role of intentions should be in an account of the epistemic conditions. Ginet thinks that having a malevolent intention is not needed for being blameworthy for a bad outcome, and this seems right. But perhaps having a benevolent intention is needed for being praiseworthy for a good outcome. It is also not clear if higher-order intentions and desires have a role to play in the epistemic conditions. In addition to his views on Frankfurt-style cases, Frankfurt famously defended an 'identification' or 'hierarchical' view of responsibility according to which, roughly, we are responsible to the extent that our higher-order desires line up with the first-order desires that move us to act. Frankfurt seems to think that this fits within the model of an actual-sequence view of responsibility because higher-order desires enter the actual sequence by overdetermining the act together with the relevant first-order desires (Frankfurt 1971, Sect. IV). But it is not clear that this is what happens. It is not clear, in particular, that higher-order desires are causally efficacious in the way Frankfurt imagines. Perhaps they need to exist just to 'protect' the causal efficacy of the first-order desires, without always being causally efficacious (if anything had got in the way between the first-order desire and the act, then the higher-order desire would have stepped

⁷ The distinction between metaphysical and epistemic conditions tends to correspond to the Aristotelian distinction between two kinds of 'excusing conditions': force and ignorance (roughly, force is the failure of the metaphysical condition and ignorance is the failure of the epistemic condition).

⁸ It is also not totally clear that the label 'epistemic' is the right one. For what is captured by the condition that the agent *should* have known that a certain outcome would result from his acts? It is not obvious that this is a purely epistemic condition. Ginet thinks that, when an agent is culpable for a harm that he should have known would follow from what he did (although he did not know that it would at the time), this is ultimately because there is something he did earlier, which he could foresee *then* would result in his being unable to foresee the harm later on. That is to say, culpable ignorance is always rooted in culpable non-ignorance. But this claim could be contested.

in by generating the intention to eliminate those obstacles, but, if the first-order desire results in the act, the higher-order desire remains causally inert). If so, perhaps the relevant higher-order desire is required for responsibility but only as an epistemic condition for responsibility, not as a link in the actual sequence.⁹

At any rate, I take it that actual-sequence views of responsibility are views that attempt to capture the metaphysical conditions for responsibility only, that is, the conditions that have to obtain for an agent to have the relevant kind of freedom or control required by responsibility. My suggestion, at least, is that this is how we should interpret them. If so, we should qualify claims (1) through (3) accordingly. For example, in the case of the supervenience thesis (claim (3), the claim that will be my main focus in what follows), we should read the thesis as the claim that any difference in *the metaphysical conditions for responsibility* has to be traced to a difference in the actual sequence, or as the claim that any difference in responsibility between agents *that are in the same epistemic situation* has to be traced back to a difference in the actual sequence. In what follows, I will assume that this is how the supervenience claim should be interpreted, and similarly for the other claims.¹⁰

⁹ Alternatively, an advocate of an actual-sequence view could understand the metaphysical condition as consisting of two parts, one of which would include the actually efficacious elements, and the other the internal mechanisms that back those elements up or help to sustain them. (In the introduction to their 2003, Widerker and McKenna suggest (n. 3) that identifying with the desires moving one to act could be a component of the metaphysical condition.) It does not really matter how exactly we demarcate the different conditions for responsibility, as long as we are able to isolate the part that actual-sequence views attempt to capture.

¹⁰ I think that this qualification helps address some important objections by Widerker to Frankfurt's programme (Widerker 2003). Consider a person who acts wrongly for selfish reasons, not because the act is wrong. Widerker claims that, in this case, the wrongness of the act is not part of the actual sequence leading to the act but it is relevant to the agent's blameworthiness. Frankfurt's response is that the wrongness still *explains* why he performed the act, even if it does not cause the act (Frankfurt 2003). But this is implausible. I would say, instead, that the wrongness of the act is relevant to the agent's blameworthiness only to the extent that the agent's knowing/believing that the act was wrong (alternatively, the fact that the agent should have known/believed that the act was wrong) is relevant to the agent's blameworthiness (for an account of blameworthiness along these lines, see Zimmerman 1988). If so, this case is not a problem for Frankfurt's programme because the agent's blameworthiness is accounted for in terms of the epistemic conditions. Another case by Widerker is this. A man decides to miss work out of laziness although (as he immediately finds out) he is sick. Widerker claims that the man is not blameworthy for missing work due to his sickness, but his sickness does not cause him to miss work. But one could agree that the agent is not blameworthy in this case and argue that the epistemic conditions are not met: the agent does not (or should not) believe that he would be acting wrongly if he were to miss work. (What if he never becomes

So far I have stated two main ways in which one could try to account for the claim that responsibility is a function of the actual sequence: claims (1) and (3). In addition, claim (2) is a consequence of claim (1). Claim (2) could thus serve as a partial characterization of actual-sequence approaches. But it cannot be considered to be a *full* characterization of such views. For the fact that mere possibilities are irrelevant to an agent's responsibility does not entail that the agent's responsibility is exclusively a function of the actual causal sequence issuing in X. This is because, even if mere possibilities were irrelevant to the agent's responsibility, actual facts *that are not links in the causal sequence issuing in X* could still be relevant to it. In other words, it is one thing to say that responsibility for X is a matter of what actually happens, and it is another thing to say that responsibility for X is a matter of the way in which X actually comes about, or of the actual sequence issuing in X. It is this last, stronger claim that is motivated by Frankfurt-style cases. Thus it is this last claim that an actual-sequence approach to responsibility (as I am understanding it) strives to capture. For Frankfurt-style cases suggest that all that matters for responsibility is how an act (omission, event, etc.) actually comes about, or what the actual explanation of that act (omission, event, etc.) is, not just that all that matters for responsibility is what actually happens.

I emphasize this because in the literature there is a tendency to characterize actual-sequence views as views according to which mere possibilities are irrelevant to responsibility (alternatively, as views according to which responsibility does not require alternative possibilities of action or the ability to do otherwise; this characterization is Fischer and Ravizza's own (see Fischer and Ravizza 1998, p. 37; see also Fischer's characterization in his 1999a, n. 56). Again, I think that these are not the best possible characterizations of actual-sequence views, or of the conception of responsibility that is motivated by Frankfurt-style cases.¹¹

aware that he is sick? In that case I would agree with McKenna (2008) that he is blameworthy for failing to go to work. After all, it seems that our epistemic situation could be such that we are blameworthy for failing to act in ways that are not required of us—for example, I can be blameworthy for failing to push a button that I think would save many lives, even if pushing the button in fact would have no such effect. McKenna thinks that the agent who misses work out of laziness is blameworthy for missing work even when he is aware that he is sick. I am less sure about this.)

¹¹ Fischer and Ravizza take themselves to be actual-sequence theorists despite their assigning some significance to unactualized possibilities. They believe that this makes sense in so far as the possibilities they fix upon are meant only to establish the dispositional nature of the elements figuring in the actual sequence (in particular, they are meant to establish

In addition to claims (1), (2), and (3), I will be interested in the following claim:

- (4) *Irrelevance of no-responsibility scenarios for responsibility scenarios:*

When an agent is responsible for X, unactualized possibilities where the agent is not responsible for X (due to a lack of control) do not help ground the agent's responsibility for X

This claim is another consequence of claim (1) (in fact, it is a consequence of claim (2) as well). Fischer has explicitly endorsed this claim and has relied on it to argue against 'flicker of freedom strategies' — strategies that attempt to show that alternative possibilities of some kind are always required for an agent to be responsible (Fischer 1994b).¹² Fischer thinks that there are strong reasons to believe in claim (4), including reasons that go beyond the general reasons to believe in the irrelevance of unactualized possibilities. This is because, Fischer argues, possibilities could only matter for responsibility (if at all) to the extent that they are sufficiently 'robust.' But when the only existing possibilities are pathways along which the agent is not responsible because he is not in control (for example, a scenario where the mad scientist intervenes and forces the agent to act in a certain way), the possibilities are not sufficiently robust.

Fischer argues for this in the following way. Suppose you believe that an agent is not responsible for an act when there is just one open pathway into the future (no alternative possibilities at all), because the agent then lacks the required sort of control. Now add just one alternative possibility: an alternative possibility along which the agent does something differently, but then ends up performing the same act and, again, lacks responsibility for it, say, because he is manipulated by a mad scientist. How could the fact that such an alternative is open to the agent transform a scenario of no responsibility into a scenario of responsibility? It would be very mysterious to suggest that it could

the 'responsiveness' of the actual mechanisms). But it is hard to say for sure which, if any, of claims (1) through (3) Fischer and Ravizza would endorse. I take it that they would reject (1) and (2) (although they might embrace a revised version of (1) according to which only links in the actual sequence *or properties that those links have*—including perhaps modal properties—can ground an agent's responsibility). On the other hand, Fischer has explicitly endorsed the supervenience claim for acts (see Fischer 1987, p. 72). But it is less clear that he would endorse the supervenience claim for omissions or for outcomes.

¹² See also Fischer 1999a, Sects II.E and II.F, and Pereboom 2001, pp. 23–8.

(it would seem to involve a kind of ‘alchemy’; Fischer 2007). For, if the agent lacked the required sort of control in the first scenario, the addition of alternative possibilities where he lacks control could not possibly increase the level of control that he actually has. In other words, the existence of alternative possibilities where an agent is not responsible because he is not in control cannot be part of what makes an agent responsible in the actual scenario.¹³

I have distinguished four different claims that could be used to give either full or partial characterizations of actual-sequence approaches to responsibility. In the sections that follow I examine the plausibility of each of these claims in light of apparent counterexamples.

2.

I will start by discussing the plausibility of claims (1), (2), and (3) with respect to responsibility for outcomes in the world. I will return to the other forms of responsibility (responsibility for acts and omissions) and to claim (4) in later sections.

I believe that claims (1) and (2) run into fatal problems. When applied to outcomes, these claims read:

- (1) *Irrelevance of causally inefficacious factors (outcomes):*
When an agent is responsible for an outcome, only links in the actual causal sequence issuing in the outcome ground the agent’s responsibility for that outcome
- (2) *Irrelevance of mere possibilities (outcomes):*
When an agent is responsible for an outcome, unactualized possibilities do not ground the agent’s responsibility for the outcome

To see the kind of problems that these claims run into, consider the following scenario (or pair of scenarios):

SWITCH A train is running out of control down a track. The train approaches a switch. Flipping the switch would send the train down a side track for a while, but the tracks reconverge up ahead, before

¹³ At this point Fischer draws an analogy with knowledge. He claims: the existence of possible scenarios where the agent ‘gets it wrong’ cannot explain the agent’s having knowledge that *p* in the actual scenario. If anything, the possible scenarios that could explain the agent’s actual knowledge that *p* are scenarios where the agent’s beliefs line up with the actual states of the world in the right way. (Fischer 1994b, p. 46.)

the location of a victim who has been tied to the track. An agent, who is standing by the switch, has reason to believe that a large fragment of track is missing from the main line between the switch and the victim's location. So he has reason to believe that, unless he flips the switch, the train will continue on the main track and derail, and thus the victim will survive. Given that he wants the victim to die, he flips the switch. As a result, the train travels on the side track, then back on the main track again after the tracks come back together, and ends up killing the victim.

Distinguish two scenarios of this kind:

SWITCH 1 A fragment of the main track between the switch and the victim's location was indeed missing.

SWITCH 2 The main track was intact (unbeknownst to the agent, someone had reconnected it earlier that morning).

The victim dies in both cases. Now, it seems that whereas the agent is responsible for the victim's death in SWITCH 1, he is not responsible for the death in SWITCH 2. In SWITCH 2, the fact that the victim would have died even if he had failed to flip the switch seems to relieve the agent of responsibility for the death (he is still responsible for acting with a bad intention, but not for the victim's death).¹⁴

This scenario suggests that claim (1) is false. For consider the state of the main track between the switch and the victim's location (in SWITCH 1 the track is missing a piece and in SWITCH 2 it is not). In neither case is the state of that segment of track a link in the *actual causal sequence* of events leading to the victim's death, for in neither case does the train run on that part of the track. But, in SWITCH 1, where the agent is responsible for the death, the state of the main track helps ground his responsibility, since he would not have been responsible for the death if the track had been intact, as in SWITCH 2. This suggests that, when an agent is responsible for an outcome, factors that are not links in the causal sequence issuing in

¹⁴ I discuss this case in Sartorio 2005. Scenarios that bear a structural resemblance to SWITCH 2 have been discussed in the literature. A famous case is van Inwagen's runaway horse case (van Inwagen 1983, pp. 176–7): Ryder is riding Dobbin, a runaway horse. He cannot slow him down but he can direct him in the ways he wants. When they are approaching a certain crossroad, Ryder recognizes one of the roads as a road to Rome. He steers Dobbin into the road he knows leads to Rome motivated by the hope that the runaway horse will kill one of the citizens. Unbeknownst to Ryder, all roads lead to Rome. Van Inwagen claims (and I agree) that Ryder is not responsible for the fact that the runaway horse passes through Rome. See also the scenarios discussed by Rowe (1989) and Fischer and Ravizza (1998, pp. 94–5).

the outcome can ground the agent's responsibility for the outcome. In other words, claim (1) is false.

For similar reasons, SWITCH suggests that claim (2) is false. In SWITCH 1, the unactualized possibility consisting in the fact that the victim would have survived if the agent had failed to flip the switch grounds the agent's responsibility for the outcome. For the agent would not have been responsible otherwise (if the train would have still killed the victim had the agent not flipped the switch, as in SWITCH 2).

How about claim (3)? When applied to outcomes, (3) reads:

(3) *Supervenience (outcomes):*

There cannot be a difference in responsibility for an outcome without a difference in the actual causal sequence issuing in the outcome

At first sight, it seems that SWITCH undermines this claim too. For, apparently, there is a difference in responsibility between SWITCH 1 and SWITCH 2 without a difference in the causal sequence. The causal sequence seems to be exactly the same in both cases: the train hurtles down the track, the agent flips the switch, the train turns onto the side track, then back on the main track again, and runs over the victim. I will argue, however, that there is a plausible reading of (3) on which (3) is not refuted by SWITCH. Supervenience can be preserved even if the other claims fail.

The main idea is to suggest that the causal *sequence* is not the same thing as the causal *history* of the outcome. The causal history of an outcome consists of the direct causes of the outcome, the causes of those causes, and so on up to some point far back enough in the chain. But if, as some people (myself included) believe, causation is not transitive, then the causes of the causes of X might not themselves be causes of X. My suggestion is that we should take the causal sequence to be those events in the causal history of an outcome that are, in fact, causes of the outcome. This narrower interpretation of 'the causal sequence' is well motivated. For, presumably, if an event is part of the causal history of an outcome without causally contributing to the outcome itself, then the occurrence of that event is not relevant to the agent's responsibility *for that outcome* (although, of course, it could be relevant to the agent's responsibility for other things).

So consider, again, SWITCH 1 and SWITCH 2. If you are one of those people who reject the transitivity of causation, it is very natural

to say that there is a failure of transitivity in SWITCH 2; indeed, cases with this structure are usually presented as potential counterexamples to the transitivity of causation.¹⁵ For in that case flipping the switch causes the train to run on the side track, and the train's running on the side track causes the victim's death, but, intuitively, flipping the switch does not cause the victim's death (this is by contrast with SWITCH 1, where flipping the switch does cause the victim's death).

Now, the question naturally arises: In virtue of what is there such a failure of transitivity in SWITCH 2? Why is flipping the switch not a cause of the victim's death in that case? An obvious part of the explanation is the fact that the victim would have died anyway if the switch had not been flipped. But this is not the full explanation. For there are scenarios (notably, 'causal pre-emption' scenarios) where an event causes an outcome even if the outcome would still have occurred in the absence of that event. For example, an assassin can cause the death of his victim by shooting him even if the victim would still have died in the hands of a backup assassin if he had failed to shoot him himself. Here the assassin pre-empts the intervention by the backup assassin and causes the victim's death, even if the death would still have occurred if the assassin had not shot him. Now, those who reject the transitivity of causation on the basis of switching scenarios would want to draw a contrast between 'switches' and 'pre-emptors.' The intuitive difference between them is that, whereas switches merely deflect pre-existing threats, pre-emptors launch new threats or otherwise help to sustain pre-existing threats. For example, diverting an ongoing train from one working track to another is deflecting or redirecting an existing threat, but shooting at a victim is creating or starting a threat for the victim. It is in virtue of this difference, it seems, that pre-emptors are causes and switches are not. Obviously, this is all at a very intuitive level. But this is not the place to try to make this distinction more precise. The intuitive characterization of the distinction will have to do for our purposes (for attempts to make it more precise, see Yablo 2000, Sartorio 2005, and Hall 2006).

I conclude that, despite initial appearances, the causal sequence leading to the victim's death is not the same in both cases: in one case it includes the flipping of the switch but in the other case it does not. As a result, claim (3) is not undermined by SWITCH.

¹⁵ Several people have noted the intuition of intransitivity in those cases, including people who prefer to stick with the transitivity thesis and explain away the intuition (such as Hall (2000) and Paul (2000)) and people who argue that we should try to capture it in our theories of causation (such as Yablo (2002), Sartorio (2005), and the later Hall (2006)).

The difference in the agent's responsibility between SWITCH 1 and SWITCH 2 can be traced back to a difference in the causal sequence issuing in the death. Claims (1) and (2), on the other hand, cannot be salvaged in this way. There is no denying that events that are not actually causes of an outcome sometimes ground the agent's responsibility for the outcome. The state of the main track is relevant to the agent's responsibility in SWITCH 1 even if it does not cause the victim's death. Similarly for unactualized possibilities: what would have happened if the agent had not flipped the switch is relevant to the agent's responsibility in SWITCH 1, even if the agent actually flipped the switch.

What this shows is that claim (3) is in an important way weaker than (and consequently more plausible than) claims (1) and (2). Even if responsibility for an outcome supervenes on the actual causes of the outcome, events that are not causes of an outcome as well as unactualized possibilities can account for an agent's responsibility for the outcome because they can determine which events *are* causes of the outcome. For example, the state of the main track in SWITCH 1 (the fact that it was disconnected) accounts for the fact that the agent's flipping the switch causes the victim's death, and it is in virtue of this that it is relevant to the agent's responsibility for the death. Similarly for unactualized possibilities: what would have happened if the agent had not flipped the switch in SWITCH 1 (the fact that the victim would have survived in that case) is relevant to the agent's responsibility for the victim's death because it accounts for the fact that the agent's flipping the switch caused the death in that case.

In the next section I will extend the results of this section to the other forms of responsibility: responsibility for acts and omissions. Note that I still have not discussed the plausibility of claim (4). I will return to that claim in section 4.

3.

Now let us consider responsibility for omissions. There are many examples in the literature that make trouble for claims (1) and (2) in the case of responsibility for omissions. Here is one such scenario:

CHILD An agent is walking along the seashore when he sees a child drowning and crying for help. He has good reason to believe that he

could easily rescue the child. Unmoved by the situation, he decides not to intervene and the child drowns.

Distinguish two scenarios of this kind:

CHILD 1 The agent could have easily rescued the child if he had tried.

CHILD 2 Unbeknownst to the agent, the agent could not have rescued the child. If he had tried to rescue the child, some hungry sharks that were swimming in the water would have attacked him and prevented him from saving the child.¹⁶

The agent fails to save the child in both cases. But it seems that, whereas the agent is responsible for failing to save the child in CHILD 1, he is not responsible for failing to save the child in CHILD 2 (he is responsible for not trying to save him, but not for failing to save him). Whether there were sharks in the water that would have thwarted the rescue attempt seems relevant to whether the agent is responsible for his failure to save the child.

This suggests that claim (1) is false, when applied to omissions. In CHILD 1, where the agent is responsible for his failure to save the child, the fact that the water was shark-free helps account for his responsibility. For, had there been sharks in the water, as in CHILD 2, he would not have been responsible for his failure to save the child. But the water's being shark-free is clearly not a link in the actual sequence issuing in the agent's omission in CHILD 1. Hence, when an agent is responsible for an omission, factors that are not links in the actual sequence of events leading to his omission can still ground his responsibility. In other words, claim (1) is false for omissions. Similarly for claim (2): what would have happened if the agent had jumped into the water to try to rescue the child (in particular, the fact that he would not have been attacked by sharks) helps account for the agent's responsibility in CHILD 1, even if it is a mere possibility.¹⁷

¹⁶ This is a case by Fischer and Ravizza (1991, p. 261, and 1998, p. 125). See also van Inwagen 1983 (Sect. 5.4), Ginet 2003, and Widerker 2003 for similar examples.

¹⁷ Frankfurt (2006) tries to explain the difference between scenarios like CHILD 1 and CHILD 2 by claiming that there simply is no omission in a case like CHILD 2. However, even if Frankfurt were right that in CHILD 2 there is no omission in one sense of 'omission', there could be different concepts of omission and the agent could still omit to act according to a broader concept of omission. In fact, when they presented the case, Fischer and Ravizza explicitly embraced such a broader concept of omission.

Again, however, claim (3) remains unscathed on the interpretation I have suggested. If we take the actual sequence leading to an omission to be the events in the causal history of the omission that are causes of the omission itself, then supervenience still holds.¹⁸ For consider: What causes the agent's failure to save the child in CHILD 1 and in CHILD 2? In CHILD 1, it is plausible to say the following: the agent deliberates about what to do, forms the intention not to jump into the water, fails to jump into the water, and as a result fails to save the child. All of the links in this sequence are, on the face of it, causes of the agent's failure to save the child. But not so in CHILD 2. Arguably, although the agent also fails to jump into the water and fails to save the child in CHILD 2, his not jumping into the water does not cause his not saving the child and thus is not part of the actual sequence in that scenario. For, given the presence of the sharks, he would not have been able to save the child even if he had jumped into the water. Compare: I fail to push the brakes and an accident happens. But the brakes were malfunctioning in such a way that the accident would not have been prevented by my pushing the brakes. Presumably, my failing to push the brakes is not a cause of the accident. Similarly for the agent's failure to save the child in CHILD 2: if the death of the child would not have been prevented by his jumping in, his failure to jump in is not a cause of his failure to save the child.¹⁹

Again, then, the difference in responsibility between CHILD 1 and CHILD 2 can be traced back to a difference in the causal sequence. The agent's failure to jump in is a link in the causal sequence leading to the agent's omission in CHILD 1 but not in CHILD 2. As a result, whereas the agent in CHILD 1 is responsible for his failure to save the child, the agent in CHILD 2 is not. The agent in CHILD 2 is responsible for his failure to jump in, but his responsibility does not extend

¹⁸ Here, as in other places in this paper, I am assuming that it is possible for omissions to be caused, and to be causes. Without this assumption the actual-sequence view would have to be revised to accommodate responsibility for omissions (presumably, one could do this by appeal to a non-causal concept of consequence, such as the concept of 'quasi-causation' developed by Dowe (2000)).

¹⁹ Note that here, unlike in the pre-emption cases discussed in the previous section, the inevitability of the outcome seems sufficient to render the omission causally inefficacious. As we have seen, intuitively pre-emptors are causes because, even if the outcome would still have occurred in their absence, they start a process that issues in the outcome. But note that the agent's failure to jump into the water in CHILD 2 does not start a process that issues in his failure to save the child, and my failure to push the brakes does not start a process that issues in the accident in the case of the malfunctioning brakes. For more on the topic of the causal powers of omissions, see Sartorio 2011.

to the failure to save the child because his failure to jump in does not cause his failure to save the child. This suggests, again, that unactualized possibilities and causally inefficacious factors can help account for an agent's responsibility for an omission because they can determine which events are causes of the omission, even if they are not causes themselves. Responsibility can still be a function of the actual sequence, in the sense captured by the supervenience claim.

Finally, let us turn to responsibility for acts. I believe that, by mimicking the structure of SWITCH (the scenario we discussed in connection with responsibility for outcomes), it is possible to come up with counterexamples to claims (1) and (2) in the case of responsibility for acts. Here is one such scenario:

SERUM An agent wakes up and finds himself hooked up to a machine that, he learns, is circulating a serum intravenously in his body. He is told that the serum induces violent acts in people (acts that they cannot resist performing). A tied-up victim (a famous violinist) is brought into the room. Before enough of the serum has reached his brain, the agent learns that the serum can take one of two possible routes to the brain and that, although he cannot unhook himself from the machine, he can switch from the route for which it is currently set to the alternative route. He has reason to believe that the preset route is blocked. Still, out of envy of the famous violinist (having tried to pursue a music career himself and having failed miserably), the agent decides to switch to the alternative route. When enough of the serum reaches his brain, the agent (under the effect of the serum) kills the violinist with his bare hands.

Distinguish two scenarios of this kind:

SERUM 1 The preset route was in fact blocked. Had the agent not switched to the alternative route, he would not have killed the violinist.

SERUM 2 The preset route was not blocked. Had the agent not switched to the alternative route, he would still have killed the violinist.²⁰

SERUM has a switching structure just like SWITCH. Hence, if we think that the agent is not responsible for the outcome in

²⁰ One could also imagine a similar scenario where the switching structure takes place inside the agent's brain (say, a scenario where the agent can redirect certain electrical impulses in his brain from a preset route to an alternative route).

SWITCH 2, we should probably also think that the agent is not responsible for his act in SERUM 2. And this is, I think, as it should be. For it seems to me that, whereas the agent is responsible for his act of killing the violinist in SERUM 1, he is not responsible for that act in SERUM 2. In SERUM 2, the fact that he would not have been able to stop the serum from reaching his brain by failing to switch to the alternative route seems to relieve the agent of responsibility for his act of killing the violinist (he is still responsible for his bad intentions, and for not trying to prevent the killing, but not for the act of killing itself).

If so, claims (1) and (2) are false for acts too. For the state of the preset route in SERUM 1 (the fact that it was blocked) helps account for the agent's responsibility in SERUM 1, even though it is not a link in the actual sequence leading to the agent's act of killing the violinist. Hence, once again, factors that are not links in the actual sequence can account for an agent's responsibility for an act. Similarly for unactualized possibilities: what would have happened if the agent had not switched to the alternative route (the fact that the serum would not have reached his brain) helps account for the fact that he is responsible for his act of killing the violinist. Possibilities can ground an agent's responsibility for an act, even if they are not actualized.

Again, we can preserve supervenience by understanding the concept of the actual sequence in the way I have suggested. For, presumably, the agent's switching to the alternative route causes his act of killing the violinist in SERUM 1 (where the preset route was blocked) but it does not in SERUM 2 (where the preset route was not blocked), for the same reason that the agent's flipping the switch causes the victim's death when the fragment of main track is disconnected but it does not when it is intact in SWITCH. Hence the difference in responsibility between SERUM 1 and SERUM 2 can be traced back to a difference in the causal sequence leading to the act in each case. Unactualized possibilities and causally inefficacious factors can help account for an agent's responsibility because they can determine which factors are causes of the agent's act, even if they themselves are not causes of the act.

I conclude that the idea that responsibility (for acts, omissions, and outcomes) is a function of the actual sequence can be preserved (at least, it can survive the main objections) when it is understood in the right way. Understanding it in the right way requires understanding it as a supervenience claim—roughly, the claim that any difference in responsibility is due to a difference in the actual causes.

The supervenience claim is consistent with the relevance of causally inefficacious factors and of unactualized possibilities to responsibility. These factors can be relevant to responsibility to the extent that they can determine what the actual causes are.

4.

Finally, let us turn to claim (4). When applied to outcomes, (4) says:

- (4) *Irrelevance of no-responsibility scenarios for responsibility scenarios (outcomes):*

When an agent is responsible for an outcome, unactualized possibilities where the agent is not responsible for that outcome (due to a lack of control) do not help ground the agent's responsibility for the outcome

In what follows I will suggest that claim (4) is false: an agent's lack of responsibility and control in a merely counterfactual scenario can help ground the agent's responsibility in the actual scenario. This, I will claim, is consistent with an actual-sequence approach to responsibility, when this kind of approach is captured by the supervenience claim (3) (in the way I have suggested we should understand this claim).

Consider the following scenario:

ASSASSINS 1 Smith wants Jones dead. Smith knows about the presence of two assassins in town. He knows that the first assassin's violent tendencies are triggered when he is exposed to a very bright shade of red. The second assassin, on the other hand, does not need to be triggered: he is determined to shoot Jones as soon as he gets his chance (given the way the assassins are positioned, this happens to be a few seconds after the first assassin gets his chance). Smith wants to be part of the killing of Jones, so, on the occasion where Jones approaches the two assassins, Smith puts on a bright red shirt that triggers the first assassin to shoot. Jones dies as a result.

It seems to me that Smith is responsible for Jones's death in this case. He could have 'opted out' of the killing of Jones but did not; instead, he willingly participated in the process that ended in Jones's death by putting on his red shirt and triggering the first assassin. Thus he is responsible for Jones's death. Compare: I could contribute to the killing of someone by providing an assassin with the bullets that he will use to kill the victim. If I do so, I am responsible for the victim's death. This is so even if someone else would have shot the victim

himself if the assassin had not been able to do so because his gun was not loaded. In particular, there seems to be an important difference between these scenarios and, again, ‘switching’ scenarios where the agent does not causally contribute to the outcome because, intuitively, all he does is redirect an existing threat (and in a way that does not change the outcome). Intuitively, putting on the red shirt that triggers the assassin’s shooting or providing the assassin with the bullets that he needs to kill his victim is not just redirecting an existing threat but, at the very least, helping to sustain an existing threat.

Now, I submit that Smith’s responsibility for Jones’s death in ASSASSINS 1 is partly grounded in the fact that he would not have been responsible for the death (due to a lack of control) if he had failed to put on the red shirt. In the counterfactual scenario where Smith fails to put on the red shirt, Jones still dies, at the hands of the second assassin. Smith is not in control of what happens to Jones and thus is not responsible for Jones’s death in this scenario (I am assuming that he could not have prevented the second assassin from shooting Jones). I will argue that his not being responsible in the counterfactual scenario due to his lack of control partly grounds his being responsible in the actual scenario.

To see this, contrast ASSASSINS 1 with a slightly different case:

ASSASSINS 2 Again, the two assassins are in town. This time, there is a way to stop the second assassin. But what stops the second assassin is also what triggers the first: it is, again, being exposed to a bright shade of red (imagine that, when in the presence of that bright shade of red, the second assassin is incapacitated — say, he is knocked unconscious — for a few minutes). Smith knows all of this. He decides to put on the red shirt, and Jones is shot dead by the first assassin.

Here it seems wrong to hold Smith responsible for Jones’s death. For there is nothing he could have done that would have resulted in his being less of a ‘contributor’ in the situation. By putting on the red shirt, he brought about the shooting by the first assassin. But he would have brought about the shooting by the *second* assassin if he had failed to put on the red shirt. He would have brought this about by omission (that is, in virtue of his failing to put on the red shirt).²¹ This is

²¹ Recall that we are assuming that omissions can be causes (see n. 18).

an important difference between ASSASSINS 1 and ASSASSINS 2. In ASSASSINS 1 the second assassin would have shot independently, not triggered by what Smith did. So, if Smith decides not to put on the red shirt, he is clearly not involved in the course of events leading to the death. But in ASSASSINS 2 there is a sense in which he is 'involved' regardless of what he does. He is 'involved' in the sense that he causes the events leading to the death regardless of what he does. In cases like this, where the agent would not have been less of a contributor by acting differently, it seems unfair to hold the agent responsible for the outcome in virtue of the way he acted. So this suggests that, in cases like ASSASSINS 1, where the agent *would* have been less of a contributor (or not at all a contributor) by acting differently, the fact that he would have been less of a contributor (or not at all a contributor) if he had acted differently helps ground the fact that the agent is actually responsible for the outcome.

When I first introduced claim (4), I noted that there is a persuasive argument in its favour (Fischer's argument, discussed in Sect. 1). Fischer's main thought was that alternative possibilities where the agent is not responsible due to a lack of control in the alternative scenario cannot provide the agent with any kind of control that is relevant for his responsibility in the actual scenario. He thought that it would be outright mysterious to suggest that possibilities where the agent lacks responsibility because she lacks control can provide the relevant kind of control for how the agent acts in the actual scenario. But it seems that there is, after all, a kind of control that these counterfactual scenarios can provide and that is relevant for responsibility. Consider ASSASSINS 1 again. The fact that, if Smith had failed to put on the red shirt, he would not have contributed to the killing, and thus the fact that he would not have been responsible in that case due to a lack of control, is, in part, why we think that Smith had the relevant kind of control in virtue of which he is actually responsible for Jones's death. Although Jones would have died regardless of what Smith had done, Smith's contribution to the killing of Jones would not have been the same regardless of what he had done. And this seems relevant to his responsibility.²²

²² In his 1999b, Fischer considers and rejects the suggestion that the ability to avoid responsibility can ground responsibility. As I note below, although I disagree with Fischer's claim that the absence of responsibility cannot ground responsibility, I believe that the view that results is consistent with an actual-sequence view of responsibility.

Our discussion of ASSASSINS 1 and ASSASSINS 2 suggests the following principle:

- (5) *Relevance of no-responsibility scenarios for responsibility scenarios (outcomes):*

If an agent is responsible for an outcome in virtue of A-ing, then, had the agent failed to A, he would not have been responsible for the same outcome by failing to A. Moreover, the agent's not being responsible for the outcome by failing to A in the counterfactual scenario grounds his responsibility for the outcome in the actual scenario.²³

In other words: no-responsibility scenarios are relevant to responsibility scenarios in that acting in a certain way does not make an agent responsible for an outcome unless failing to act in that way would not have made the agent responsible for that outcome.²⁴ Moreover, the fact that acting in a certain way makes the agent responsible for the outcome is partly grounded in the fact that failing to act in that way would not have made the agent responsible for it. For example, if wearing the red shirt makes Smith responsible for Jones's death, then this is partly in virtue of the fact that Smith would not have been responsible for the death if he had not worn the red shirt. Similarly, if flipping a switch and redirecting the runaway train towards a side track makes you responsible for a death that ensues, then this is partly in virtue of the fact that not flipping the switch would not have made you responsible for the same death. Also, if deciding on your own to shoot someone makes you responsible for your victim's death, then this is partly in virtue of the fact that not deciding on your own to shoot him (say, deciding to do this only as a result of being manipulated by a mad scientist) would not have made you responsible for his death. And so forth. In general, whenever an agent is responsible for something in virtue of having acted in a certain way, this is

²³ To this we should probably add: if an agent is responsible for an outcome in virtue of *failing to A*, then, had the agent A-ed, he would not have been responsible for the outcome by A-ing. Moreover, his not being responsible for the outcome by A-ing in the counterfactual scenario grounds his responsibility for the outcome in the actual scenario.

²⁴ Note that the claim is not that the agent would not have been responsible for the outcome if he had failed to act in the relevant way, but, rather, that he would not have been responsible for the outcome *by virtue of failing to act in that way*. Imagine that Smith was responsible for the presence of the two assassins in town. In that case he would presumably be responsible for Jones's death; still, he would not be responsible for the death by virtue of having failed to put on his red shirt. Thanks to an anonymous referee for raising this point.

partly in virtue of the fact that the agent would not have been responsible by failing to act in that way, if he had failed to act in that way.

As it stands, claim (5) only concerns responsibility for outcomes. It would be easy to generalize it to instances of responsibility for acts/omissions where an agent is responsible for the act/omission in virtue of something else he did/failed to do. But it is not clear how to make it more general than that.²⁵ Never mind. It is enough for my purposes to show that, at least for some kinds of responsibility (responsibility for outcomes, or responsibility that an agent bears for something in virtue of having acted/failed to act in a certain way), no-responsibility scenarios can help ground the agent's responsibility.

Certain philosophers who have defended 'flickers of freedom' strategies have put forth claims that on the surface appear to be similar to claim (5). For example, Wyma argued for a principle according to which, roughly, an agent is responsible only if he could have avoided responsibility ('principle of possibly passing the buck'; Wyma 1997). Similarly, McKenna argued that an agent is responsible only if he could have avoided 'authorship' of his action (McKenna 1997). And Otsuka argued for a principle according to which, roughly, an agent is blameworthy only if he could have behaved in a manner for which he would have been entirely blameless ('principle of avoidable blame'; Otsuka 1998). But these principles are in fact quite different from claim (5). For they are 'flickers of freedom' principles: they are designed to show that responsibility requires alternative possibilities of some kind. Roughly, they claim that responsibility requires *the ability to avoid responsibility*, even if it might not require the ability to do otherwise. Claim (5), on the other hand, does not entail that responsibility requires alternative possibilities of any kind. In particular, it does not entail that responsibility requires the ability to avoid responsibility. (5) only says that responsibility requires that the agent would have avoided responsibility *if he had done otherwise*. From this it does not follow that the agent had the ability to avoid responsibility, unless the agent had the ability to do otherwise. As a result, someone who opposes the alternative-possibilities requirement for

²⁵ Suppose that we try the following: if the agent is responsible for X in virtue of Y (Y could be something else he did/failed to do, or some relevant part of the causal chain leading to X), then, if Y had not occurred/had been different in some way, the agent would not have been responsible for X in virtue of Y's not occurring/being different in that way. This is an implausible generalization. There are many different ways in which the causal chain leading to X could have been different, and the agent could still be responsible for X due to those things being different in many of those scenarios.

responsibility would have no problem embracing claim (5). In fact, an opponent of the alternative-possibilities requirement for responsibility could help himself to claim (5) to explain whatever initial appeal flicker-of-freedom views like Wyma's, McKenna's, and Otsuka's have. He could say that claim (5) captures the 'grain of truth' in those views: the important but *only* truth behind them. So when I claimed that I disagree with Fischer about his claim that the absence of responsibility or control cannot ground responsibility, I did not mean to suggest that I disagree with him about his rejection of the alternative-possibilities requirement on responsibility. I think that one can accept that (perhaps surprisingly) responsibility is grounded in the absence of responsibility or control, without believing that responsibility requires alternative possibilities of any kind.

I will end by arguing that embracing claim (5) is, in fact, fully consistent with embracing an actual-sequence approach to responsibility. In particular, I will argue that it is consistent with embracing claim (3) (the supervenience claim), as I have proposed that we should understand this claim. To see this, consider, again, ASSASSINS 1 and ASSASSINS 2. Drawing on the earlier discussion, we can say that there is a difference in the causal sequence between the two scenarios, which grounds the difference in moral responsibility between them. In particular, we can say that Smith's wearing the red shirt causes Jones's death in ASSASSINS 1 but it does not in ASSASSINS 2. In ASSASSINS 2, there is a failure of transitivity: Smith's wearing the red shirt causes the first assassin to shoot, which causes Jones's death, but Smith's wearing the red shirt does not cause Jones's death. Putting on the red shirt is like flipping the switch that diverts the train from the main track onto the side track: it may or may not be a cause of the final outcome, depending on what would have happened if one had acted differently. And, if Smith's putting on the red shirt causes Jones's death in one scenario but not in the other, then this is an important difference that can ground the moral difference between the two cases. For it means that Smith's contribution to Jones's death is different in the two scenarios.²⁶

²⁶ In Sartorio 2005 I argued that causes are 'difference-makers' in the sense that, if an event causes an outcome, then its absence would not have caused the same outcome. In ASSASSINS 1, Smith's putting on the red shirt causes Jones's death, and his failing to put on the red shirt would not have caused Jones's death. However, in ASSASSINS 2, if we were to say that Smith's putting on the red shirt caused Jones's death, then there would be similar pressure to say that his not putting on the red shirt would also have caused it. But this is ruled

My final conclusion is, then, this. Not only is an actual-sequence view of responsibility consistent with the claim that unactualized possibilities are *sometimes* relevant to responsibility; it is also consistent with the claim that unactualized possibilities are *always* relevant to responsibility. In fact, I have argued that the most plausible actual-sequence view is one that holds that unactualized possibilities of a certain kind are always relevant to responsibility.²⁷

References

- Aguilar, J., K. Frankish, and A. Buckareff (eds) 2011: *New Waves in the Philosophy of Action*. Basingstoke: Palgrave-Macmillan.
- Dowe, Phil 2000: *Physical Causation*. Cambridge: Cambridge University Press.
- Fischer, John Martin 1987: 'Responsiveness and Moral Responsibility'. In Fischer 2006a, pp. 63–83. Originally published in Schoeman 1987.
- 1994a: *The Metaphysics of Free Will*. Oxford: Blackwell.
- 1994b: 'Responsibility and Alternative Possibilities'. In Widerker and McKenna 2003, pp. 27–52. Originally published in Fischer 1994a.
- 1999a: 'Recent Work on Moral Responsibility'. *Ethics*, 110, pp. 93–139.
- 1999b: 'Responsibility and Self-Expression'. *The Journal of Ethics*, 3, pp. 277–97.
- 2006a: *My Way: Essays on Moral Responsibility*. New York: Oxford University Press.
- 2006b: 'A Framework for Moral Responsibility'. In Fischer 2006a, pp. 1–37.
- 2007: 'Response to Kane, Pereboom, and Vargas'. In Fischer, Kane, Pereboom, and Vargas 2007, pp. 184–90.
- Fischer, John and Mark Ravizza 1991: 'Responsibility and Inevitability'. *Ethics*, 101, pp. 258–78.

out by the difference-making principle. I think that, if I am right and the difference-making principle about causation is true, this helps explain why (5) is true.

²⁷ Many thanks to Juan Comesaña, John Martin Fischer, Michael McKenna, Derk Pereboom, an audience at the University of California–Los Angeles, the members of a graduate seminar on moral responsibility at the University of Wisconsin–Madison in the spring of 2009, and the Editor and four referees for *Mind*.

- 1998: *Responsibility and Control*. Cambridge: Cambridge University Press.
- Fischer, J., R. Kane, D. Pereboom, and M. Vargas (eds) 2007: *Four Views on Free Will*. Oxford: Wiley-Blackwell.
- Frankfurt, Harry 1969: 'Alternate Possibilities and Moral Responsibility'. *Journal of Philosophy*, 66, pp. 829–39.
- 1971: 'Freedom of the Will and the Concept of a Person'. *Journal of Philosophy*, 68, pp. 5–20.
- 1994: 'An Alleged Asymmetry Between Actions and Omissions'. *Ethics*, 104, pp. 620–3.
- 2003: 'Some Thoughts Concerning PAP'. In Widerker and McKenna 2003, pp. 339–45.
- Ginet, Carl 2000: 'The Epistemic Requirements for Moral Responsibility'. *Philosophical Perspectives*, 14, pp. 267–77.
- 2003: 'In Defense of the Principle of Alternative Possibilities: Why I Don't Find Frankfurt's Argument Convincing'. In Widerker and McKenna 2003, pp. 75–90.
- Hall, Ned 2000: 'Causation and the Price of Transitivity'. *Journal of Philosophy*, 97, pp. 198–222.
- 2006: 'Structural Equations and Causation'. *Philosophical Studies*, 132, pp. 109–36.
- McKenna, Michael 1997: 'Alternative Possibilities and the Failure of the Counterexample Strategy'. *Journal of Social Philosophy*, 28, pp. 71–85.
- 2008: 'Frankfurt's Argument Against Alternative Possibilities: Looking Beyond the Examples'. *Noûs*, 42, pp. 770–93.
- Otsuka, Michael 1998: 'Incompatibilism and the Avoidability of Blame'. *Ethics*, 108, pp. 685–701.
- Paul, L. A. 2000: 'Aspect Causation'. *Journal of Philosophy*, 97, pp. 223–34.
- Pereboom, Derk 2001: *Living Without Free Will*. Cambridge: Cambridge University Press.
- Rowe, William 1989: 'Causing and Being Responsible for What is Inevitable'. *American Philosophical Quarterly*, 26, pp. 153–9.
- Sartorio, Carolina 2005: 'Causes as Difference-Makers'. *Philosophical Studies*, 123, pp. 71–96.
- 2011: 'Failing to Do the Impossible'. In Aguilar, Frankish, and Buckareff 2011, pp. 32–49.
- Schiffer, Stephen 1987: *Remnants of Meaning*. Cambridge, MA: MIT Press.

- Schoeman, Ferdinand (ed.) 1987: *Responsibility, Character, and the Emotions: New Essays on Moral Psychology*. Cambridge: Cambridge University Press.
- van Inwagen, Peter 1983: *An Essay on Free Will*. New York: Oxford University Press.
- Watson, Gary (ed.) 2007: *Free Will*. New York: Oxford University Press.
- Widerker, David 2003: 'Blameworthiness and Frankfurt's Argument Against the Principle of Alternative Possibilities'. In Widerker and McKenna 2003, pp. 53–73.
- Widerker, David and Michael McKenna (eds) 2003: *Moral Responsibility and Alternative Possibilities*. Aldershot: Ashgate, rpt 2006.
- Wyma, Keith 1997: 'Moral Responsibility and Leeway for Action'. *American Philosophical Quarterly*, 34, pp. 57–70.
- Yablo, Stephen 2002: 'De Facto Dependence'. *Journal of Philosophy*, 99, pp. 130–48.
- Zimmerman, Michael 1988: *An Essay on Moral Responsibility*. Totowa, NJ: Rowman and Littlefield.