

Juggling Intuitions about Causation and Omissions

Carolina Sartorio*

1 Introduction

Intuitions are central to philosophical theorizing about causation, but there are different kinds of intuitions, which can play different kinds of roles. In this chapter, I focus on the following four types:

Causal intuitions: Intuitions to the effect that something is/isn't a cause of a given outcome.

Explanatory intuitions: Intuitions to the effect that something is/isn't part of the explanation of something else—where the relevant notion of explanation tracks something potentially broader than just causation.

Responsibility intuitions: Intuitions to the effect that agents are/aren't morally responsible for a given outcome.

Grounding intuitions: Intuitions to the effect that agents are/aren't morally responsible for a given outcome *because*, or *to the extent that*, they are/aren't a cause of (or part of the explanation of) that outcome. These intuitions track grounding relations between causation (or explanation) and moral responsibility.

I focus on these four types of intuitions because I believe that they are particularly relevant to philosophical theorizing about causation. (I don't mean to suggest that these are the *only* sets of intuitions to be deserving of consideration, but just that they are some central ones.) In particular, my main focus here will be, not these intuitions taken in isolation from each other, but the *interplay* between them. I'll be examining questions such as these: In what ways do these different intuitions interact with one another? How should these interactions inform our theorizing about causation? What should we do when there's a conflict between intuitions of different kinds? And so on. The main examples I'll work with are cases of *omission* (failures to act). The reason for this is that, as we will see, in cases of omission the interaction between intuitions of

different kinds is particularly significant for the purposes of theorizing about causation. Hence, they make for an interesting case study.

Causation theorists typically focus on intuitions of the first kind, *causal intuitions*, insofar as these can be directly used to motivate philosophical theories of causation or to critically evaluate them. However, intuitions of the other kinds are also relevant to philosophical theorizing about causation, in a more indirect but still important way. Causation is not an isolated concept but one that is connected with other theoretically useful concepts, including, in particular, the concept of explanation and the concept of moral responsibility.¹ As a result, it can be illuminating to look at the broader picture that includes these other concepts and the interrelations among them. As we will see, this is particularly important in cases where our causal intuitions themselves aren't particularly clear, such as cases of omission. (Another example of this that I will discuss in this chapter is cases of causal overdetermination.)

Interestingly, this perspective is rarely taken in the literature on causation. This chapter is an attempt to remedy this. I believe that an investigation of the concept of cause won't be exhaustive unless we think about the way causation fits with those other key concepts. Accordingly, one of the main goals in this chapter is to raise awareness about the centrality of this issue.

Some preliminaries are in order before we start.

First of all, I must note that, as a metaphysician, the concept of cause which I'm interested in is a *natural* concept. That is, it's a concept that picks out a certain relation between events or states of affairs, one that exists "out there in the world." This is a concept that is importantly connected to some normative concepts, such as the concept of moral responsibility, but it's not itself a normative concept. In particular, being causally responsible doesn't require being morally responsible for something. Natural events such as tornadoes are causally responsible for outcomes, but they are not morally responsible for anything. The same goes for moral agents like us: we can be causally responsible for outcomes without being in any way morally responsible for them, as when we trip over someone accidentally and faultlessly cause them harm.

Second, I see intuitions the way I think many other theorists do, as starting points or as data that should be taken into account in our theorizing, but almost never as the last word. In my view, intuitions about causation are important to the extent that they help us latch on to the relevant relation in the world (the one that we're trying to pick out with our concepts and language). However, given, in particular, the connections that exist between causation and other concepts such as responsibility and explanation, it can be hard to know how to "juggle" all the different kinds of intuitions at once, and how to strike the best balance among them. Such a process of *reflective equilibrium* can result in our paying more attention to some intuitions rather than to others, and even to jettison some altogether, when we formulate our theories.²

Third, here I won't be relying on empirical studies, but mostly on what I consider to be "educated guesses" about commonly shared intuitions about causation and the other connected concepts. My belief is that most of these educated guesses would quite reliably track lay people's intuitions. Others may not, at least not as reliably—but not because they would necessarily clash with them, but simply because raising the issues in an intelligible matter requires a bit more philosophical sophistication or training. I

think it could be interesting, and to some extent illuminating, to run empirical studies on some of these educated guesses, especially those involving a certain kind of scenario of omission that I'll focus on later, but I'll leave that task to others.

More generally, and despite my mainly theoretical focus, I hope this chapter will be of interest to those working on empirical debates about our concept of causation, given, in particular, the new questions it raises about the interactions among the different kinds of intuitions, and the special challenges to which those interactions give rise. I believe that paying attention to those unexplored challenges can be fruitful when theorizing about causation, both from a purely theoretical perspective and from an empirically informed one.

Again, I predict that the majority of readers will share my intuitions about the cases presented in this chapter, but even if you don't, some of those same challenges are still likely to arise. I'll walk you what I take to be the most common intuitions about cases, and the particular questions that those intuitions provoke. Even if you don't share some of these intuitions, similar questions may arise for you as you attempt to juggle your own intuitions. We all tend to have intuitions of the four kinds above, and it isn't always clear what the best way to accommodate them is.

2 Introducing cases of omission and overdetermination

I'll start with some relatively "easy" cases—cases where causation theorists are mostly in agreement about the role played by the relevant intuitions. We will then consider more complex scenarios in which there is more disagreement: cases of omission and cases of overdetermination.

As mentioned above, causation theorists typically focus on intuitions of the first kind, *causal intuitions*, to formulate or confirm their theories. Sometimes these intuitions are so pervasive and powerful that they can settle important theoretical questions pretty much on their own. These "bedrock" intuitions are the basis on which much theorizing about causation is done. A good example of this is intuitions about **preemption cases**. For instance, philosophers typically focus on scenarios of this kind:

Fast and Slow: Two agents, Fast and Slow, throw rocks at a fragile and valuable vase. Fast's rock reaches its target and breaks the vase right before Slow's rock, which sails through empty space. If Fast's rock hadn't broken the vase, the vase would have broken anyway, and in a very similar way, as a result of Slow's rock hitting it.³

What caused the vase to shatter? The answer seems obvious: it was Fast's throw, not Slow's throw, but it's notoriously hard to accommodate this simple fact within a general theory of causation. In particular, some popular views that attempt to analyze causation in terms of the notion of counterfactual dependence face the challenge of explaining why it is that Fast's throw caused the vase's shattering, when the shattering doesn't counterfactually depend on Fast's throw.⁴ Much philosophical work in this tradition has acknowledged the force of these bedrock causal intuitions, and has focused on finding solutions to problems of this kind.⁵

Incidentally, note that cases like **Fast and Slow** can also be used to illustrate the force of some *grounding intuitions*. Imagine that we know that only one of the rocks hit the vase, but we don't know which one that was. Imagine, also, that both agents acted intentionally, freely, etc. In that case, we can still be in a position to know that whoever broke the vase is *morally responsible for its shattering*, and the other agent isn't. This is based on a grounding intuition: the intuition that an agent's moral responsibility for the broken vase is grounded, among other things, in having caused that outcome. As a result, the agent who didn't cause the outcome cannot be morally responsible for it⁶ (although, of course, she can still be responsible for other things, such as for trying to break the vase, for acting with a malevolent intention, and so on). For the purposes of this chapter (given our focus on causation and the causal upshots of our behavior), this is the type of responsibility that we're mainly interested in: responsibility for *outcomes*, or a moral assessment of agents in light of the upshots of their behavior. When we ask about who is morally responsible for the outcome in this case, it's only natural to look at who is causally responsible.

Of course, in this example there are certain empirical facts of which we are unaware, and this is what results in the uncertainty about the causes of the outcome. That uncertainty can be eliminated simply by coming to know the relevant empirical facts (whose rock hit the target on that occasion). This could easily lead to the impression that, to the extent that we are aware of all the relevant empirical facts about particular cases, the corresponding causal intuitions will always be sufficiently clear. However, this is actually far from the truth, for causal intuitions can fail to be fully clear or universal *even in cases where we know all the relevant empirical facts*. These are the kinds of cases that are of particular interest to us here, because they point to a genuine unclarity about how to theorize about causation itself (assuming the empirical facts are settled). Not coincidentally, they are also the kinds of cases for which the interaction between the different types of intuitions mentioned above becomes particularly relevant. In what follows, I'll draw attention to two examples of this phenomenon.

The first example, which will be our main focus here, is cases of **omission** (or absences in general) and the lively philosophical debate concerning whether omissions can be causes.⁷ This debate is typically fueled by a more basic debate about the nature of the *causal relata* (the "terms" of the causal relation, or the kinds of things that the causal relation relates). Some think that only "positive" things like ordinary events can be causal relata, and this seems to rule out omissions and other absences; others, however, disagree. In the context of this debate, we could know all the relevant facts about agents' omissions (including everything that agents haven't done but could have done, and all the facts concerning what would have happened if they had done those things, etc.), and this still wouldn't come close to settling the basic philosophical issue of whether omissions can be causes, or what the causal relata in general are.

Of course, there are some powerful intuitions to the effect that agents *can* cause outcomes by omission. For example, philosophers have focused on scenarios of this kind:

Dead plants: I hired a gardener who committed to caring for the plants in my backyard. He failed to tend to my plants (e.g., he didn't water them), and my plants died. They would have lived otherwise.

In this case, it seems very natural to regard the gardener as a cause of the plants' death. In response, however, some argue that intuitions about causation involving omissions are still not as clear or forceful as intuitions about causation involving positive events. If a thunderstorm uproots all the plants in my backyard, the thunderstorm seems to be more clearly a cause of my plants' demise than any omission by the gardener that would have prevented that outcome (imagine, for example, that the plants would have been spared if he had protected them with thick cloth coverings). Dowe (2001) calls this "the intuition of difference." He then goes on to argue that omissions lack causal efficacy, partly on the basis of the intuition of difference.⁸

Another potential reason to discount the significance of causal intuitions involving omissions, such as the intuition concerning **Dead plants**, is that it's notoriously difficult, or even impossible, to accommodate all those causal intuitions within a general theory of causation. For intuitions about causation by omission tend to be infused by normative considerations that are in tension with the aspirations of a philosophical account of causation conceived of as a *natural* concept (and this concept, recall, is my focus here).⁹ As a result, no general account of this concept is likely to be able to capture the whole range of causal intuitions involving omissions.

We can illustrate this point with the same example from before. Notice that, although the gardener appears to be a cause of my plants' death in **Dead plants**, Tucson's mayor (Regina Romero) doesn't, but we can imagine that Regina Romero bears all the same natural relations to my plants' death as my gardener (in particular, she could have dropped by my home in Tucson and watered my plants, and the plants would have survived if she had watered them). The main difference is that it was *the gardener's job* to water them (and not Regina Romero's); hence, only the gardener is morally responsible for the plants' death. In other words, our responses in these cases seem to be tracking, at bottom, *moral responsibility intuitions*. However, again these judgments about moral responsibility are tracking normative considerations that don't bear on the natural concept of cause. As a result, an account of that concept will have to ignore some of the causal intuitions about omissions.

Finally, imagine that, as some philosophers believe, omissions cannot be causes. Still, it is surely possible to account for the significance of the gardener's omission in other terms, without appealing to causation. On this alternative "fallback" view, the gardener's failure to care for my plants is *part of what explains* why they died, even if it's not a cause of the plants' death. Notice that this is what's captured by the third kind of intuitions mentioned above: *explanatory intuitions*. Many authors who reject the causal efficacy of omissions in fact embrace the idea that omissions can be part of the full explanation of events (see for example, Dowe, 2001; Beebe, 2004; and Varzi, 2007). This idea can be put to use to account for the agents' moral responsibility in those cases, in accordance with the corresponding *grounding intuitions* (by claiming that moral responsibility is grounded in, if not the causal powers of omissions, at least their explanatory power).

What could be meant by "explanatory power," you may ask, if not a causal power? Without giving a precise account of this concept, the idea is that, if omissions and other absences cannot be causes, they can still contribute to the full explanation of events in that those events still happen, at least partly, *because* of those absences. For example,

the plant died, at least partly, because it wasn't watered. (After all, if somebody had watered it, it wouldn't have died.) It is also common to suggest—and the authors mentioned above do suggest—that, if omissions cannot be causes, we can still capture the explanatory power of omissions in terms of a form of causation. This time it's not *actual* causation, but *counterfactual* causation, however. That is, we can say that omissions help explain events, not because of the causal relations that actually obtain, but because of the causal relations that *would have* obtained if the omissions hadn't taken place. In terms of possible worlds, this is the idea that omissions are explanatorily powerful by virtue of causal relations that obtain in possible worlds relevantly similar to the actual world. For example, the non-watering of a plant can help to explain why it died in that, in close possible worlds where the plant is watered, it survives.

So far, we have identified cases of omission as one main set of examples where causal intuitions aren't fully clear or pervasive, even when we are aware of all the relevant empirical facts. In those cases, as we have seen, there is a lack of clarity about causation itself. We also noted the relevance, in those cases, of the intuitions of the other types (responsibility intuitions, explanatory intuitions, and grounding intuitions) and the relations among them. As we have seen, looking at the causal intuitions as part of a larger net of intuitions can help us get a better perspective on things, from which we can see the different options that open up when theorizing about causation.

Another example of the same type of phenomenon is **symmetric overdetermination cases**. These cases have also been the subject of lively philosophical debates. I'll go through these a bit more quickly.

Philosophers have illustrated the phenomenon of symmetric overdetermination with examples of the following kind:

Two Rocks: Imagine that two agents throw rocks at a vase, but this time the rocks hit the vase simultaneously. Imagine, also, that each rock would have been sufficient on its own for the vase to break (in roughly the same way, and at roughly the same time).¹⁰

An important difference between a symmetric overdetermination scenario like **Two Rocks** and the **Fast and Slow** scenario described above (a preemption case) is that in **Two Rocks** the potential causes (the two rocks, or the two rock-throwing events) are fully on a par—hence the label “symmetric.” Thus, it is not the case that one of them is a cause while the other one isn't, or that one “preempts” the other.

Plus, although it's clear that the vase broke, somehow, thanks to the two throwers' actions, this doesn't fully settle how we should think about the causal structure of the case. For consider: Should we say that each of the rocks or rock-throwing events was an individual cause of the vase breaking? (Some argue for this position. But note that this results in more causes than is needed to explain the effect, which some people find objectionable.) Or should we say, instead, that the cause is a fact that is more “proportionate” to the effect—perhaps the fact that *somebody* threw a rock at the vase? Or some “collective” event? (However, what kinds of facts or events are these? How can they be causes without the individual events being causes?) Again, the precise causal

structure of a scenario like **Two Rocks** remains unclear even if we know all the relevant empirical facts about the case.¹¹

Notice that, here too, we have an explanatory claim to fall back on, in case we conclude that the best way to make sense of the causal structure of the case is to say that overdeterminers aren't individual causes. For, surely, even if overdeterminers aren't individual causes, they *help explain* the outcome, in some important sense, considering the collective contribution they make to it. In other words, the full explanation of the vase's shattering will have to appeal, in some way or other, to the contribution made by the two rocks. The vase didn't just break for no reason!

In any case, here I'll understand the term "explanatory" in this very broad way, as an umbrella term that captures contributions of different kinds—including, in particular, some non-causal contributions as well as some collective contributions. If we're looking for a real "fallback" option to explain what happened, and to potentially ground the moral responsibility of the agents involved, this broad notion seems to be the best candidate.

To sum up: in this section, we have identified two types of case where causal intuitions tend to be particularly less clear or pervasive than in other (more ordinary) cases, and where this is due to a genuine uncertainty about causation, or about the conceptual tools needed to make sense of certain causal scenarios (and not about the underlying empirical facts). In the next section, I look more closely at the interplay of intuitions of different kinds that takes place in those cases, and I offer a diagnosis.

3 The resilience of moral responsibility and the flexibility of grounding intuitions

The scenarios discussed in the previous section can help bring out the *flexibility of grounding intuitions* in our theorizing about causation. So far, we have alluded to purely *causal* grounding intuitions, when discussing the **Fast and Slow** preemption scenario. In that case, I pointed out, it's very natural to take moral responsibility to be straightforwardly grounded in, among other things, causation. On the basis of that causal grounding intuition (and other things, such as the fact that the agents in question were acting freely, with a bad intention, etc.), we tend to conclude that whoever broke the vase (the preemptor) is also, thereby, morally responsible for the vase's breaking. This is so even if we might not know who that is, or who preempted whom, but must grounding intuitions *always* play this same kind of role? Must they all be *causal* grounding intuitions?

Presumably not. For consider, again, cases of omission and cases of symmetric overdetermination. Imagine that we come to believe that our best theories of causation imply that omissions and overdeterminers are never causes. Would we be tempted to conclude, on that basis, that the agents in those cases simply lack any moral responsibility for the outcome (because they didn't cause those outcomes)? Again, presumably not. Instead, we'd be willing to *relax* the relevant grounding claim, in a way that would allow for the agents' moral responsibility in those cases to be grounded in something other than individual causal relations.

In the previous section we noted that, in both kinds of cases, we have some explanatory intuitions to fall back on. In **Dead plants**, we could still blame the gardener for the death of my plants, for his omission would still be part of the explanation of my plants' death (in the broad sense described above) even if it were not a cause. In **Two Rocks**, we could still blame the two agents who threw the rocks at the vase and broke it simultaneously, because their behavior would still be part of the explanation of the vase's breaking (again, in the broad sense described above), even if they were not individual causes of that outcome. I take it that most of us would be ready to appeal to that fallback option, rather than letting intuitively culpable agents off the hook.

To clarify, when I say that agents seem clearly responsible in these cases and that it would be a mistake to let them off the hook, I simply mean that it's clear that they bear *some* moral responsibility for the outcome. This is consistent with claiming, for example, that agents are less blameworthy in virtue of their omissions than by virtue of their actions.¹² It's also consistent with claiming that agents involved in symmetric overdetermination cases are less blameworthy than if they had been the only agents involved.¹³ I won't take a stand on these issues. All I'm interested in here is the claim that the agents would not simply be off the hook, in that they would still bear some moral responsibility for what happens. I take this to be a non-negotiable intuition.

Now, why is it that not letting the agents off the hook seems like the right reaction to have about these cases? I propose the following diagnosis. Judgments about moral responsibility are, to some important extent, resistant to certain philosophical—in particular, metaphysical—discoveries, such as discoveries about the true nature of causation. This includes the discovery that omissions or overdeterminers aren't causes.

Let's give this idea a label:

The resilience of moral responsibility: Many judgments about responsibility are “resilient” in that they would survive certain metaphysical discoveries about the nature of causation.¹⁴

My suggestion, then, is that the resilience of moral responsibility is what results in the flexibility of the corresponding grounding claim. The flexibility of the grounding claim is the idea that we (most of us, anyway) are prepared to, if needed, relax the purely causal grounding intuition and instead rely on a substitute or surrogate intuition of the following kind:

Relaxed grounding intuition: Moral responsibility for outcomes is grounded in, if not causation, then, more broadly, *explanation* (in the broad sense of explanation described above).

The readiness to switch from the purely causal version to this relaxed version of the grounding claim seems to strike the best balance between the different types of intuitions that we have about these cases. For it allows us to preserve resilient judgments about moral responsibility while holding on to a close analog of the causal grounding claim: the claim that moral responsibility is grounded in, more broadly than just causation, explanatory power.

I would in fact push for a slightly revised (and, I believe, improved) version of the relaxed grounding claim. This is the claim that moral responsibility for outcomes is grounded in, a bit more precisely, *moral responsibility* for some explanatory factors. I believe that this modification helps account for cases where some agents contribute to the full explanation of an outcome, but where they are still not responsible for the outcome because they are not responsible for the explanatory factors themselves. I have argued for this in earlier work and I won't go into any of the details here.¹⁵ For present purposes, we can sidestep this complication for the most part, so I'll only make reference to it when needed.

Let me take stock of what we have so far. I started by distinguishing four different kinds of intuitions that can contribute to our theorizing about causation (in potentially different ways): causal intuitions, explanatory intuitions, responsibility intuitions, and grounding intuitions. As part of the discussion of *causal intuitions*, I described cases where those intuitions are so clear and pervasive that they tend to act as bedrock intuitions (my example of this was preemption cases). I then contrasted these scenarios with cases where the causal intuitions are much less clear (my examples were omission cases and symmetric overdetermination cases), and I discussed the interaction that takes place in those cases between intuitions of the four different kinds. I argued that, in those cases, where the judgments about moral responsibility tend to be quite resilient, and thus where the *responsibility intuitions* are particularly strong, the relevant *explanatory intuitions* act as a fallback resource that can be put to use (if needed) in accounting for the agents' moral responsibility. This, in turn, can be accomplished by relaxing the causal grounding claim, in a way that respects the substance of the relevant *grounding intuitions*.

In the next section, I will consider more complex scenarios involving omissions that raise special and more difficult challenges. In these cases, as we will see, the interplay between the different types of intuitions seems to come apart from what I have described in this section.

4 A special challenge: asymmetric omission cases

The case of omission discussed above, **Dead plants**, is a "simple" omission case where an agent fails to do something he was supposed to do and the outcome happened, apparently, as a result of that omission. But there are other omission cases that are much more complex. These cases can resist an easy treatment. In this section, I explain how I see the interaction between the intuitions of different types in these interesting cases.

One way to think about a *slightly* more complex omission case is to incorporate the features of the other type of scenario discussed above, the one involving symmetric overdetermination. This yields a **symmetric overdetermination case** involving omissions. Consider, for example, the following scenario:

Symmetric flooded room: Some valuable art pieces are kept inside a room. They are located in an area where heavy rains are common. When it rains heavily, the

flooding of the room can be prevented by simultaneously closing *two doors*, door 1 and door 2 (a single door is not enough to stop the rainwater's flow). Two agents, A and B, are in charge of operating those doors (one door each, because the switches are in different locations). When the alarm sounds at time T, letting the agents know that they must pull their switches, both of them simultaneously fail to do this, purely out of laziness (neither agent is aware of what the other agent intended to do at T). The art pieces are ruined.¹⁶

This case is an omission case which is otherwise similar to **Two Rocks**, our earlier example of symmetric overdetermination. Since both doors needed to be shut to prevent the flood, each agent's omission is *sufficient by itself* to guarantee the occurrence of the flood and the destruction of the art pieces. Plus, the agents' contributions are perfectly symmetric or on a par (in particular, the relevant omissions are failures to close the doors *at exactly the same time*), so this means that there cannot be preemption of one omission by the other. For these reasons, this is a symmetric overdetermination case, albeit one involving omissions.

Now, in this case, the responsibility judgments seem to be just as resilient as those about **Two Rocks**: each agent seems to bear at least some moral responsibility for what happened. After all, the flood wouldn't have occurred had it not been for the behavior of the two agents, and that behavior was blameworthy (each of them should have done their part, and they had no good excuse for not doing so—in particular, they had no reason to believe that the other agent would *also* fail to do their part).¹⁷ Notice also that, here too, we can use the relaxed grounding claim and the relevant explanatory claim to account for the agents' responsibility, if needed (if omissions or overdeterminers aren't causes). For we can say that what each of them did (or, in this case, *failed* to do) is part of the full explanation of what happened, and this explanatory role can be used to ground their responsibility.

Let's now see what happens when we turn it into a different kind of case: an **asymmetric overdetermination case**. This is a case where one of the two omissions precedes the other. That is, the relevant omissions are failures to behave in certain ways at *different* times (an earlier time and a later time).¹⁸ For example, consider the following asymmetric variant of the flooded room case:

Asymmetric flooded room: The setup is similar to that of the symmetric case, except that door 1 can only be closed at time T1 and, door 2, at a later time T2. Imagine that A is supposed to close door 1 at T1, and B is supposed to close door 2 at T2. Again, neither agent is aware of what the other agent intends to do, since they are in separate rooms. And, again, imagine that both agents independently fail to do their job, purely out of laziness. As a result, they each independently fail to close their doors, and the art pieces are ruined.¹⁹

This type of case raises some unique problems.

There is still *a* judgment about moral responsibility that seems resilient in this case, but it's not the claim that both agents are responsible. Rather, it's the claim that *somebody* is responsible (either A or B, or both). Somebody must be responsible, for, again, the art

pieces wouldn't have been ruined if it hadn't been for the blameworthy behavior of two agents. So, we can't just excuse both agents. This is, arguably, a non-negotiable intuition. But what's interesting about this case is that now it's no longer perfectly clear *who* is responsible: if one, or the other, or both. (In symmetric cases, it was clear that it was both, in light of the perfect symmetry of the case.) As a result, the more specific judgments about moral responsibility are less clear in this case: we're not as sure who to blame.

Why is this, exactly? One main reason is that, given that the case is no longer perfectly symmetric, preemption reappears as a live option: given the temporal asymmetry, A could be preempting B, or B could be preempting A (or the equivalent of that for non-causal explanatory relevance, if omissions turned out not to be causes: A could be explanatorily relevant instead of B, or B instead of A). Or it could be that, despite the temporal asymmetry, this is still a case of symmetric overdetermination (or, again, the equivalent of that for non-causal explanatory relevance: both are equally explanatorily relevant).²⁰

This uncertainty exists even though we know all the relevant empirical facts. This means that it's an uncertainty about the causal or explanatory power of omissions; in particular, it's an uncertainty about the conditions under which causal or explanatory preemption happens, *for omissions*. For notice that cases involving ordinary ("positive") events and causal connections don't raise equally difficult challenges. In those cases, we have physical causal processes to look at. If two rocks are thrown at a vase, for example, we can follow the rocks' paths and we know that there is preemption if one rocks get to the target and breaks it first. But in omission cases there are no such physical processes, but *absences of physical processes*. As a result, there is nothing to trace.

Also, asymmetric cases of overdetermination involving omissions are especially challenging because in these cases one could potentially argue that the two behaviors "cancel each other out"—that is, one could argue that the first omission isn't causally or explanatorily relevant because the second omission renders it irrelevant, and vice-versa. Thus, in **Asymmetric flooded room**, one could argue as follows: "B was going to fail to close door 2 later (at T2) regardless. So, the fact that A failed to close door 1 *earlier*, at T1, is simply irrelevant: closing that door wouldn't have made any difference to the outcome." However, in a parallel fashion, one could argue as follows: "A had *already* failed to close door 1 (at T1). So, the fact that B failed to close door 2 *later*, at T2, is simply irrelevant: closing that door wouldn't have made any difference to the outcome." Of course, when we combine these two pieces of reasoning, we're led to the conclusion that neither behavior was causally or explanatorily relevant. However, this is unbelievable: on the basis of this, and the relevant grounding intuition, we would be able to conclude that neither agent is morally responsible for what happened, but as noted above, this seems unacceptable. So, something went wrong in this reasoning, but it's hard to say what it is.

Finally, another thing that makes these cases puzzling is that the responsibility judgments are likely to change significantly when we imagine variants on the cases where one of the agents is replaced by a mechanism or a non-agential, natural phenomenon. Imagine, for example, that there is no agent B, and that in its place door 2 was going to be closed automatically at time 2. Imagine, however, that the mechanism

fails, or that lightning strikes at that precise moment and destroys it before it's activated. In that case, doesn't agent A seem less responsible for the art pieces being ruined? After all, they were bound to be ruined as a result of an "act of God" (the legal terminology that is typically used for this kind of thing), and not as a result of two agents failing to do their job. But, how can we explain the difference in responsibility between this "natural" variant on the case and the original version? The contribution that agent A herself makes seems to be the same in both cases. In both cases, A fails to close door 1 when door 2 was not going to be closed for independent reasons; whether those reasons concern another agent or a natural phenomenon is arguably irrelevant to the causal contribution that A herself makes.²¹

So, let me return to what I think we do know for sure about a case like **Asymmetric flooded room**: we know that someone is responsible for what happened. What we do not know is who is responsible. Now, what could possibly determine who is responsible? At this point, it seems that knowing what the *actual* causal or explanatory structure of the case is could help. For it would tell us whose behavior was *in fact* relevant to what happened.

In other work, I have explained what I take the causal (explanatory) structure to be, and the implications for the responsibility of the agents involved. The answer isn't simple, but it's not important for our purposes in this chapter.²² Here I'm only interested in the general issue concerning the special kind of interplay that takes place between the different types of intuitions in these cases. I think it's clear that it's quite unusual; in particular, it seems to be different from what we have seen in the previous section.

How does the interaction between intuitions work, in these cases? What general lessons can we learn from this?

One thing that these cases seem to show is that the **resilience of moral responsibility**, despite being an important and widespread phenomenon, may only be *limited*. That is to say, sometimes we can be *genuinely uncertain* about who is responsible and who is not, even when we know all the relevant empirical facts. The answer in those cases seems to hinge on who is a cause, or who is explanatorily relevant (or—I would say—on who is morally responsible for the cause or explanatorily relevant factors), and we must do some heavy-duty philosophical work to figure this out.²³

This is unlike what happens in the simpler cases, where we know who is responsible (assuming we know all the relevant empirical facts), and the philosophy we need to do consists in figuring out the best way to conceptualize this—if in terms of causation, or a non-causal form of explanatory relevance. With the asymmetric cases, in contrast, we must start from scratch, in a way. For we have to figure out, *at once*, and presumably by means of a delicate exercise of reflective equilibrium, who is responsible and who is explanatorily relevant (or responsible for the explanatorily relevant factors). As a result, given that we have to do everything at once, it's hard to know where to start.²⁴

Let me end by commenting on the potential value of doing empirical research on these cases. To the best of my knowledge, there are no empirical studies on cases of this kind.²⁵ This is a shame. It would be interesting to know what people's intuitions are about these cases, but it's important to realize that these scenarios are challenging, not just for purely theoretical purposes, but also for the purposes of running empirical studies. For their structure is quite complex, in that they combine special features of

different kinds: they are *omission* cases; plus, they involve *overdetermination*; plus, they are *asymmetric*. This combination of features is what results in their being such an interesting case study, but it's also what results in inevitable complexities at the time of surveying people's intuitions.

I think it would be interesting to know, in particular, what people think about the agents' *moral responsibility* in these kinds of cases: Am I right in thinking that we're not ready to let both agents off the hook for what happened? Would people tend to blame both agents, or just one of them? If only one of them, then which one?

Here it's important to recall the important distinction between being responsible for one's *behavior* and being responsible for the *outcomes* of one's behavior. The outcome, not the behavior, is what's overdetermined in scenarios like **Asymmetric flooded room**. Thus, what we're interested in finding out about these cases is who is morally responsible for the outcome (the flood, or the art pieces being ruined), not who is morally responsible for their own behavior. However, this is a distinction that it might be easy for the ordinary person to overlook, or to underestimate, when thinking about complex cases like this. Sometimes we just want to blame blameworthy people, and we don't pay close attention to whether we're blaming them for what they did or for the results of what they did. After all, this is a subtle philosophical distinction that might take some time getting used to. So, this is a challenge that would have to be overcome in running these studies.

It would also be interesting to know what people think about the *causal or explanatory structure* of these kinds of cases: Who caused the flood? Or, whose behavior was explanatorily relevant to the flood's occurrence?

Here it's important to bear in mind the phenomenon alluded to before, in section 2: ordinary causal judgments about omissions tend to be normatively loaded. (This might be true more generally too, but, as noted above, it's *especially* true in the case of omissions.) Due to this effect, it's natural to expect that people's causal or explanatory judgments will tend to go hand in hand with their moral responsibility judgments. However, these results would have to be taken with a grain of salt, if what we're interested in is the natural concept of cause that is the main focus of metaphysical investigations.

To sum up: in this section, we have considered some complex scenarios involving omissions, asymmetric overdetermination cases, which raise special and more difficult challenges. I have argued that here the interplay between the different types of intuitions is unlike what we see in more ordinary omission cases. Correspondingly, these scenarios give rise to new theoretical questions about causation, as well as to some unique challenges for doing empirical research on the subject.

5 Conclusions

In this chapter, I have discussed the interface between intuitions of different kinds as it bears on our theorizing about causation. I focused mostly on scenarios of omission as a distinctive case-study.

Omissions are interesting to philosophers for many reasons, one of which is that they raise special puzzles about the nature of causation. I have suggested that some of

those puzzles are reflected in the interaction that takes place between different *kinds* of intuitions: intuitions about causation, explanation, moral responsibility, and intuitions to the effect that moral responsibility is grounded in causation or explanatory power. As we have seen, perhaps surprisingly, that interaction isn't always clearly the same, for all omission cases; instead, it can take different forms depending on the particular type of case at issue. This makes the study of omissions (and of different kinds of scenarios involving omissions) especially challenging, but at the same time especially interesting.

Notes

- * For helpful comments, thanks to Shaun Nichols, Alex Wiegmann, Pascale Willemsen, and an anonymous reviewer for this volume.
- 1 Some argue that there is more than one concept of cause (see for example, Hitchcock, 2007). I don't want to take a stand on this issue, and the existence of more than one concept of cause is in fact compatible with everything I say in this chapter. What the chapter assumes is only that there exists *some* concept of cause (a natural concept—as I explain next) that is connected with moral responsibility in important ways. For a discussion of this assumption in the context of the problem of free will and moral responsibility, see Sartorio (2016), especially chapter 2.
- 2 For a more extended discussion of this perspective, see Paul and Hall (2013), especially pp. 2–4 and 41–2. There intuitions are described as “defeasible guides to potentially interesting and important features of our causal concept or a causal relation” (p. 2). This is how I see them too. However, Paul and Hall don't discuss the role played by the interplay between the different types of intuition I focus on in this chapter (they work mostly with causal intuitions).
- 3 This is a classical “late preemption” case (see Lewis, 1986, postscript E).
- 4 An event Y counterfactually depends on another event X when the counterfactual conditional “If X hadn't occurred, then Y wouldn't have occurred” is true (X and Y are both actual events). In **Fast and Slow**, the vase's shattering doesn't counterfactually depend on Fast's throw because it would still have occurred (as a result of Slow's throw) if Fast's throw hadn't occurred.
- 5 See for example, the papers collected in Collins, Paul, and Hall (2004).
- 6 This is what gives rise to the interesting phenomenon of resultant moral luck. I offer an analysis of this concept in Sartorio (2012a).
- 7 For example, see the debate between Dowe and Schaffer in Hitchcock (2004). See also Bernstein (2015) for a general overview.
- 8 Of course, there are other ways of accounting for the intuition of difference. For example, one could claim that causation comes in degrees and that omissions and other absences are causes to a lesser degree than positive events. (However, the view that causation is a scalar notion is quite unpopular among causation theorists—for an overview of the relevant literature, see Kaiserman, 2018. I have argued against the intelligibility of degrees of causation in Sartorio, 2020.) One could also account for the intuition of difference without appealing to a metaphysical difference between actions and omissions, but only to a pragmatic difference. For example, one could argue that omissions tend to be much less salient causes than actions, although both of them are causes.
- 9 See my comment on this in section 1 above. There has been much discussion of the influence that normative considerations have on causal judgments about omissions;

- see for example, Beebe, (2004); McGrath (2005); Livengood and Machery (2007); Clarke et al. (2015); Willemsen (2016); Henne, Pinillos, and De Brigard (2017); Henne, et al. (2019); and Willemsen (2019). For an overview of how normative considerations affect causal judgments more generally, see Willemsen and Kirfel (2018).
- 10 For a classical discussion of symmetric overdetermination cases, see Lewis (1986), postscript E.
 - 11 For a discussion of these two positions on the problem of symmetric overdetermination, see Schaffer (2003). Schaffer calls the two views “individualism” and “collectivism.” He himself defends the individualist position.
 - 12 This is usually connected with the idea that there is a moral difference between doing and allowing harm (for an overview of the doing/allowing harm debate, see Woollard and Howard-Snyder, 2016). For empirical research on asymmetries in moral judgments concerning actions and omissions, see for example, Cushman and Young (2011).
 - 13 Zimmerman (1985) discusses this view and argues against it.
 - 14 I discuss this thesis in Sartorio (2021). There I also discuss the contrast between these kinds of metaphysical discoveries and others that could potentially be relevant to moral responsibility judgments, such as finding out that determinism is true. (*Incompatibilists* about the determinism and free will problem believe that the truth of determinism would, in fact, undermine our free will and moral responsibility.)
 - 15 See Sartorio (2004, 2012b, 2015a, and 2017).
 - 16 I discussed a case with a similar structure, the “Two Buttons” case, in Sartorio (2004).
 - 17 For a dissenting opinion about this kind of case, see Moore (2009; chapters 5, 6, and 18). I respond to this aspect of Moore’s view in Sartorio (2012b). There I discuss the best interpretation of the slogan “Two wrongs don’t make a right,” and I argue that it yields the consequence that both agents are blameworthy for the outcome in this type of scenario.
 - 18 Most real-life cases, I take it, are asymmetric in this sense, given the artificiality of the symmetric cases (where, recall, the two agents needed to act at precisely the same time in order to prevent the occurrence of the outcome). Real-life cases typically involve *windows of time* during which the agents could have acted, and those windows won’t perfectly overlap. For example, imagine that Joe had part of the day to work on his contribution to a project, and Mary had another part of the day (maybe partially overlapping Joe’s); each contribution was essential to the success of the project, and both independently failed to do their part; as a result, the project failed. **Asymmetric flooded room** is a “cleaner” case, artificially designed to avoid these complications, in order to keep things as simple as possible.
 - 19 Similar cases have been discussed in the causation literature. Most of them don’t directly involve omissions, though, but actions that have some relevant absences as results, which raise similar puzzles. See for example, McDermott (1995) and Collins (2000) on “preemptive prevention” cases. See also the puzzle of the “desert traveler” case from the literature on causation in the law (see for example, McLaughlin, 1925–6, and Hart and Honore, 1985). Briefly, the puzzle is the following: A man takes a trip into the desert, carrying his water canteen. The man has two enemies, A and B, who want him to die and who independently come up with a criminal plan to make sure that happens. A first drains the water out of the canteen, and then, not noticing that the canteen is empty, B steals it. The man dies of thirst in the desert. Who killed the desert traveler?
 - 20 See Metz (forthcoming) for an argument that the first agent preempts the second in a case of this kind. (See also Zimmerman, 1985 for a similar claim about responsibility.)

- In the literature on the desert traveler puzzle (the puzzle mentioned in n. 19 above), each of the different possible positions has been defended by at least one theorist, which goes to show how much disagreement there is on this issue. (I discuss the puzzle and my own solution to it in Sartorio, 2015a.)
- 21 I argue that this type of reasoning gives rise to a new form of moral luck in Sartorio (2015b).
 - 22 Spoiler alert! My view is that A is the one who's morally responsible for the outcome, and B is not. In fact, this is not because A is explanatorily relevant and B isn't, but because A is *morally responsible* for the explanatorily relevant factors and B isn't. (See my modification of the relaxed grounding claim alluded to in the previous section.) However, I admit that not everybody would agree with this solution to the puzzle; after all, it's really hard to know what the best solution is. I discuss puzzles of this kind in Sartorio (2015a and 2017).
 - 23 Notice that, if I am right and the answer depends on who's morally responsible for the explanatorily relevant factors (as explained in section 3), then the answer will depend on another judgment about moral responsibility. But notice that such a judgment is just as uncertain: it is also not clear, at least initially, who's responsible for the explanatorily relevant factors. This is what makes the puzzle particularly hard, in my view.
 - 24 Notice that one intuition that does seem indefeasible, even in these challenging cases, is the *relaxed grounding intuition*. If one has in mind a very broad sense of explanatory relevance, as I do here, then the claim that responsibility for an outcome requires being explanatorily relevant to the outcome (in some way or other) seems simply undeniable. Personally, I can't imagine circumstances that would lead me to reject it.
 - 25 Reuter et al. (2014) contains interesting empirical research on the effect that temporal differences have on the phenomenon of causal selection. But their studies involve actions, not omissions—and joint causation, not overdetermination.

References

- Beebe, H. (2004). "Causing and nothingness." In J. Collins, N. Hall, and L.A. Paul (eds.), *Causation and Counterfactuals* (291–308). Cambridge, MA: MIT Press.
- Bernstein, S. (2015). "The metaphysics of omissions." *Philosophy Compass* 10(3): 208.
- Clarke, R., J. Shepherd, J. Stigall, R. Repko Waller, and C. Zarpentine (2015). "Causation, norms, and omissions: A study of causal judgements." *Philosophical Psychology* 28: 279–93.
- Collins, J. (2000). "Preemptive prevention." *Journal of Philosophy* 97(4): 223.
- Collins, J., N. Hall, and L.A. Paul (eds.) (2004). *Causation and Counterfactuals*. Cambridge, MA: MIT Press.
- Cushman, F. and L. Young (2011). "Patterns of moral judgment derive from nonmoral psychological representations." *Cognitive Science* 35: 1052–75.
- Dowe, P. (2001). "A counterfactual theory of prevention and 'causation' by omission." *Australasian Journal of Philosophy* 79(2): 216–26.
- Hart, H.L.A. and T. Honore (1985). *Causation in the Law*, 2nd ed. Oxford: Oxford University Press.
- Henne, P., A. Pinillos, and P. De Brigard (2017). "Cause by omission and norm: Not watering plants." *Australasian Journal of Philosophy* 95(2): 270–83.
- Henne, P., L. Niemi, A. Pinillos, F. De Brigard, and J. Knobe (2019). "A counterfactual explanation for the action effect in causal judgment." *Cognition* 190: 157–64.

- Hitchcock, C. (ed.) (2004). *Contemporary Debates in Philosophy of Science*. Malden, MA: Blackwell.
- Hitchcock, C. (2007). "Three concepts of causation." *Philosophy Compass* 2(3): 508–16.
- Kaiserman, A. (2018). "‘More of a cause’: Recent work on degrees of causation and responsibility." *Philosophy Compass* 13(7).
- Lewis, D. (1986). "Causation." In *Philosophical Papers*, Vol. II (159–213). New York: Oxford University Press.
- Livengood, J. and E. Machery (2007). "The folk probably don't think what you think they think: Experiments on causation by absence." *Midwest Studies in Philosophy XXXI*: 107–28.
- McDermott, M. (1995). "Redundant causation." *British Journal for the Philosophy of Science* 46: 523–44.
- McGrath, S. (2005). "Causation by omission: A dilemma." *Philosophical Studies* 123/1: 125–48.
- McLaughlin, J.A. (1925–6). "Proximate cause." *Harvard Law Review* 39: 149.
- Metz, J. (forthcoming). "Preemptive omissions." *Erkenntnis*.
- Moore, M. (2009). *Causation and Responsibility*. Oxford: Oxford University Press.
- Paul, L.A., N. Hall and E.J. Hall (2013). *Causation: A User's Guide*. Oxford: Oxford University Press.
- Reuter, K., L. Kirfel, R. van Riel, and L. Barlassina (2014). "The good, the bad, and the timely: How temporal order and moral judgment influence causal selection." *Frontiers in Psychology* 5(5): 1336.
- Sartorio, C. (2004). "How to be responsible for something without causing it." *Philosophical Perspectives* 18(1): 315–36.
- Sartorio, C. (2012a). "Resultant luck." *Philosophy and Phenomenological Research* 84(1): 63–86.
- Sartorio, C. (2012b). "Two wrongs don't make a right: Responsibility and overdetermination." *Legal Theory* 18: 473–90.
- Sartorio, C. (2015a). "Resultant luck and the thirsty traveler." *Method* 4(6): 153–72.
- Sartorio, C. (2015b). "A new form of moral luck?" In A. Buckareff, C. Moya, and S. Rosell (eds.), *Agency, Freedom, and Moral Responsibility* (134–49). Palgrave-Macmillan.
- Sartorio, C. (2016). *Causation and Free Will*. Oxford: Oxford University Press.
- Sartorio, C. (2017). "The puzzle(s) of Frankfurt-style omission cases." In D. Nelkin and S. Rickless (eds.), *The Ethics and Law of Omissions* (133–47). New York: Oxford University Press.
- Sartorio, C. (2020). "More of a cause?" *Journal of Applied Philosophy* 37(3): 346–63.
- Sartorio, C. (2021). "Responsibility and the metaphysics of omissions." In S. Bernstein and T. Goldschmidt (eds.), *Non-Being: New Essays on Non-Existence* (294–309). New York: Oxford University Press.
- Schaffer, J. (2003). "Overdetermining causes." *Philosophical Studies*, Vol. 114 (1/2): 23.
- Varzi, A. (2007). "Omissions and causal explanations." In F. Castellani and J. Quitterer (eds.), *Agency and Causation in the Human Sciences* (155–67). Paderborn: Mentis Verlag.
- Willemsen, P. (2016). "Omissions and expectations: a new approach to the things we failed to do." *Synthese* 195: 1587–614.
- Willemsen, P. (2019). *Omissions and their Moral Relevance: Assessing Causal and Moral Responsibility for the Things We Fail to Do*. Paderborn: Mentis Verlag.
- Willemsen, P. and L. Kirfel (2018). "Recent empirical work on the relationship between causal judgments and norms." *Philosophy Compass* 14: 1.

- Woollard, F. and F. Howard-Snyder (2016). "Doing and allowing harm." *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition), Edward N. Zalta (ed.), <https://plato.stanford.edu/archives/win2016/entries/doing-allowing/>
- Zimmerman, M. (1985). "Sharing responsibility." *American Philosophical Quarterly* 22: 115–22.