

Responsibility in Context

Ann Whittle¹

Received: 20 October 2015 / Accepted: 29 January 2017 / Published online: 2 March 2017
© The Author(s) 2017. This article is published with open access at Springerlink.com

Abstract Some have argued that our intuitive reactions to a number of cases of moral responsibility can only be preserved at the expense of a unified account of moral responsibility for acts and omissions. I argue against this conclusion, proposing that a plausible condition on responsibility, the Causal Condition can, when properly elaborated, justify the relevant intuitive data.

A theory of moral responsibility will ideally do two things. First, it will present a unified account of our responsibility for acts, omissions and their consequences. Second, it should account for as many as possible of our intuitive reactions to a wide range of relevant cases.¹ Unfortunately, there is some reason to think that these two demands pull against each other. Some have argued that our intuitive reactions to a number of cases can only be preserved at the expense of a unified account of moral responsibility for acts, omissions and their consequences. I argue against this conclusion, proposing that a plausible condition on responsibility, the Causal Condition (CC) can, when properly elaborated, justify the relevant intuitive data. So reflective equilibrium with respect to these cases can be reached, without resorting to asymmetric conditions for the consequences of acts and omissions.

In the first section I describe a number of cases, outlining the nature of the challenge for a unified account of moral responsibility. In the second section I describe CC, articulating it in terms of a contextualist account of causation. In the third section I show how, so understood, CC can explain and justify our intuitive

¹ If ‘S is morally responsible for X’ then S is deserving of some kind of reactive attitude, most notably praise or blame, because of X.

✉ Ann Whittle
ann.whittle@manchester.ac.uk

¹ School of Social Sciences, The University of Manchester, Manchester M13 9PL, UK

reactions to the cases under consideration. I end by considering an objection to the view.

1 Puzzling Cases

Consider first the following two cases:

Sharks: John is walking along a beach and sees a child drowning. He decides to ignore the child's plight. Unbeknownst to John, if he had decided to jump in, he would have been eaten alive by the large number of sharks circling the shore. (Based on Fischer and Ravizza 1998, p. 125)

Missile: Elizabeth launches a missile at a city, thereby devastating it. If she hadn't launched the missile, Carla would have done so instead, with exactly the same consequences. (Based on Fischer and Ravizza 1998, p. 94)

Most have concluded that, in *Sharks*, John is not morally responsible for the child's drowning. Of course, he is morally responsible for his woeful lack of concern for the child's plight but, since he would in any case have been unable to save the child, he is not responsible for this dreadful outcome. In *Missile*, on the other hand, it is intuitive to think that Elizabeth is morally responsible for the destruction of the city, and this is so even though there is nothing she could have done to prevent it from occurring.

This result puts some pressure on a unified account of responsibility. By 'unified' here, I mean that one set of conditions applies equally to different kinds of ascriptions of moral responsibility—be they ascriptions of moral responsibility for acts, omissions or their consequences. *Sharks* and *Missile* threatens this since, in both cases, the agent was unable to prevent the outcome in question. The only relevant difference between the two seems to be that whilst *Missile* involves an act and its consequences, *Sharks* involves an omission and its consequences. So we might be tempted to endorse an asymmetric account of moral responsibility, one on which responsibility for the consequences of omissions requires the ability to prevent the relevant consequence, but responsibility for the consequences of acts does not.²

The following two cases, however, complicate this picture:

² This position is endorsed by Sartorio, 'NA: An agent's responsibility for an action can transmit to an outcome even if the outcome would have occurred anyway in the absence of the action. However, an agent's responsibility for an omission cannot transmit to an outcome if the outcome would have occurred anyway in the absence of the omission' (2005, p. 470). Clarke (1994, 2011) and McIntyre (1994) put forward a similar thesis, the core of which Clarke refers to as '**INTAB** An agent is responsible for omitting to A only if, had the agent intended to A, he would have been able to A' (2011, p. 611). Although this is not stated in terms of the consequences of omissions, given Clarke does apply it to such cases as *Sharks*, I assume that he thinks it can so extend (2011, p. 615). So John is not morally responsible for his omission of failing to save the child.

Clarke (2014) no longer defends the asymmetry thesis, saying 'there is, I think, a decent case to be made that there are successful Frankfurt inaction cases. But I confess that the question seems to me to remain uncertain.' (p. 153) Clarke leaves it unclear how his new view fits with his earlier work (1994, 2011), as this does accommodate Frankfurt inaction cases (see Frankfurt 1994, p. 622). I have omitted a discussion of these cases here, even though they do pose a problem for asymmetric analyses, since they raise different issues which I have addressed in Whittle (2016) (see §5 for a discussion of Frankfurt-style cases involving omissions).

Two Buttons: A dangerous chemical is leaking which, if unchecked, will lead to a massive explosion. Agent A and agent B can seal off the leak by simultaneously pressing their emergency buttons. As it happens, both decide to do nothing and an explosion ensues. (Based on Sartorio, 2004, p. 317)

Stuck Button: A dangerous chemical is leaking which, if unchecked, will lead to a massive explosion. Two buttons need to be pressed simultaneously to seal off the leak. Agent C must press one emergency button; the other button should automatically be activated once a leak has been detected. Unfortunately, the button on the automated system is jammed and, in addition, agent C decides not to press her button. (Based on Sartorio 2004, p. 318)

In *Two Buttons* agent A is intuitively responsible for the explosion despite the fact that it was a consequence of an omission and she could not have prevented that consequence. If so, it is a counterexample to the view that moral responsibility for omissions and their consequences requires the ability to prevent the relevant consequence. So it seems that the simple asymmetry claim above must be rejected.³

Stuck Button, however, just adds to the puzzle since, as Sartorio plausibly claims, while we hold Agent A morally responsible for the explosion in *Two Buttons*, the same is not the case of agent C in *Stuck Button*. This is so even if we stipulate that the causal process leading up to the explosion is identical in each case. The only difference between the cases seems to be that, in *Two Buttons*, the second button's not being pressed is due to another agent's failure, whereas in *Stuck Button*, the button isn't pressed due to the failure of the automatic safety mechanism. But why should this matter?

According to Sartorio, 'the fact remains that B wasn't depressed due to a mechanical failure, not a human failure, and somehow this seems to take away my responsibility for the fact that the two buttons weren't simultaneously depressed' (2004, p. 329).⁴ But without an explanation of why it is relevant whether or not the omission is caused by an agent, this is unsatisfying. Further, there is reason to think that this is not the distinction that matters. Consider:

³ Neither Clarke's (2011) nor Sartorio's (2005) asymmetric views can accommodate *Two Buttons*. This is despite the fact that the case is borrowed from Sartorio (2004). Moreover, this kind of counterexample differs from that offered by Clarke (2014, p. 157). It cannot be dealt with by adding the following caveat 'unless the lack of such dependence is itself something for which the agent in question is responsible' (p. 157), since agent A is not morally responsible for agent B's actions, and vice versa.

⁴ A similar position is espoused by Byrd (2007). When discussing Frankfurt cases concerning omissions, he argues that the best solution to the puzzles surrounding the cases is 'to bite the bullet and conclude that natural events should not count as triggering events' (2007, p. 63). So if you couldn't have done otherwise due to something in the natural world, like in *Stuck Button*, this robs you of responsibility for the outcome. But if you couldn't have done otherwise due to the choices of other rational agents, as in *Two Buttons*, then you are not let off the hook. Byrd does not offer any explanation of why this difference, if it does hold, matters.

Full Shark: John is walking along a beach and sees a child drowning. He decides to ignore the child's plight. Unbeknownst to John, if he had decided to jump in, there is a slim possibility that he might have been eaten alive by Sharon the Shark circling the shore. Sharon has just eaten a big dinner, however, and is quite old and lazy, so John might well have been fine.

Plausibly, if the chance of the child's being killed is increased by John's inaction, he is morally responsible for it. The fact that John might have failed due to a natural rather than agential event, however, does not seem to be a pertinent excuse which would mitigate John's responsibility. More pressing is how likely it is that the rescue would have been prevented by the intervention of the shark. There is, then, reason to be doubtful of the claim that what matters is whether or not the would-be preventer is agential or non-agential.

Taken together, these cases are puzzling and put some pressure on the goal of formulating a unified account of responsibility for acts, omissions, and their consequences. First, there are omission cases in which the ability to prevent the outcome seems crucial (e.g. *Shark*) and others in which it does not (e.g. *Missile*). Further, there are omission cases in which the agential nature of the preventer seems crucial (e.g. *Two Buttons*) and others in which it does not (e.g. *Full Shark*). It may seem from all of this, then, that responsibility for the consequences of omissions poses special problems for an overall account of moral responsibility.

2 The Causal Condition

There is a familiar distinction between moral and causal responsibility. I am causally responsible, without being morally responsible, for every causal consequence of my acts and omissions. But whilst being a causal consequence of one of my acts or omissions is not a sufficient condition of moral responsibility, it is plausibly a necessary condition. That is, everything for which I am morally responsible is either an act or omission of mine, or a causal consequence of an act or omission for which I am responsible. This is the Causal Condition:

CC: If *A* is morally responsible for *e* then *e* is either an act or omission of *A*'s or is a causal consequence of an act or omission for which *A* is morally responsible.⁵

⁵ CC relies on the view that omissions can be causes. To avoid this commitment, we might rewrite CC in terms of 'quasi-causal consequences', but although I cannot argue for it here I think that such a modification is both unnecessary and undesirable. Another sort of case that may challenge CC would be cases of collective action in which causation is not distributive (Cooper (1968), Sartorio (2006 §1)) but responsibility is. Whether there are such cases is a complex question and cannot receive the treatment it deserves here. Finally, some may object to the import, if not the letter, of CC by arguing that we are never morally responsible for anything we cause (see Graham 2014 for an interesting discussion). I think that we should, in accordance with standard usage, say that people are responsible for the consequences of their actions/omissions. It is important to be able to distinguish between those bad or good events that are, at least in part, to be attributed to human agency and those which are mere 'acts of God' or terrible accidents. However, this claim is left as an assumption of the paper.

CC relies on, but does not itself provide, an account of moral responsibility for acts and omissions themselves. This account may be symmetrical between acts and omissions; for example, we may argue that the ability to do otherwise is a *ceteris paribus* necessary condition of responsibility for each of these.⁶ But although I take it that such a symmetrical account of moral responsibility is, other things being equal, to be preferred, it is not assumed. What CC offers is a necessary condition on responsibility for those events that are neither acts nor omissions of ours; we are only morally responsible for those events that are causal consequences of acts or omissions for which we are responsible.⁷

CC, I shall argue, can explain and justify our intuitive reactions to the cases discussed in the previous section. It does so, furthermore, without necessitating any asymmetry regarding the ability to prevent the outcome or any unexplained principles concerning the agential nature of preventers. But to see how, we first need to say something about causation. It is tempting for those working on moral responsibility to remain neutral on questions regarding causation. However, sometimes it is only by getting involved in issues surrounding the nature of causation that progress can be made. In particular, I shall argue that the key for making the Causal Condition work for a theory of moral responsibility is to combine it with causal contextualism.

Causal contextualism is the view that the truth-value of causal claims varies with context. So the very same causal claim might be true in one context of utterance, but false in another. This is an increasingly popular view, in some part because it is a commitment of the plausible structural equations framework.⁸ Central to all current variations on this approach is the relativisation of singular causal claims to a particular causal model. The causal model displays our conceptualization of the particular situation in which the causal claim is made. As Pearl and Halpern comment,

the truth of every claim must be evaluated relative to a particular model of the world; that is, our definition allows us to claim only that C causes E in a (particular context in a) particular structural model. It is possible to construct two closely related structural models such that C causes E in one and C does not cause E in the other. (2005, p. 845)

According to Pearl and Halpern, the structural equations represent independent causal mechanisms encoding information about how the exogenous variables (variables representing the background situation that we are taking for granted and

⁶ See Whittle (2016).

⁷ Some of our acts and omissions may be the causal consequences of antecedent acts and omissions of ours. For instance, my omission to water the flowers may be the causal consequence of my mental act of deciding not to water the flowers. If so, CC will apply to such omissions and acts. But the argument of this paper does not presuppose this claim.

⁸ At least as it is standardly construed, see, for example, Hitchcock (2001), Menzies (2004) and Woodward (2003) all of whom develop, in different ways, Pearl's structural equations framework (2000). I shall outline the basic idea to illustrate causal contextualism, but I am not committed to a structural equations framework. Schaffer (2012) offers a very useful summary of some of the data that supports causal contextualism, independent of the structural equations framework. For other, distinctive causal contextualist accounts that could be adopted see, for instance, Schaffer (2005, 2012) and Hall (2007).

whose values, which we hold fixed, are determined by factors outside the model (Halpern and Pearl 2005, p.847, 850)) determine the values of the endogenous variables (variables whose value is determined by factors within the model).⁹ By intervening and changing the value of some variable in one of these equations, whilst holding the other factors fixed, we can determine whether that variable makes a difference to the effect.

To illustrate and support the basic approach, consider this example, borrowed from Hitchcock (1996) and Menzies (2007),

Patient: A patient needs at least 100 mg of a drug to survive. He can be given 3 different doses of the drug: 0, 100, or 200 mg. In this case he is given 100 mg.

Is his taking of the 100 mg dose a cause of his survival? On the structural equations framework, to answer this we need to provide a causal model of the situation. Following Menzies (2007, p. 204), let's suppose that we can simplify it as follows:

Variables of Model:

C can take values 0mg, 100 mg and 200 mg.

R = 1 if the patient survives, 0 if not.

Structural Equations of Model:

$C = 100 \text{ mg}$

$R = f(c)$, where $f(c) = 1$ if $C \geq 100 \text{ mg}$, otherwise 0

To determine whether or not taking the drug made a difference, we give C a new value whilst leaving the other factors unchanged. So we give C either 0 or 200 mg and then recalculate the value of R. If the value of C at 100 mg makes a positive difference to whether or not the patient survives, then we can say that it is a cause. But, clearly, there is a problem. If we contrast $C = 100 \text{ mg}$ with $C = 0 \text{ mg}$, then it does make a difference as the value of R changes from 1 to 0, but not if we contrast $C = 100 \text{ mg}$ with $C = 200 \text{ mg}$.

Hitchcock and Menzies both conclude that whether consuming 100 mg of the drug caused the patient's survival is relative to some alternative or causal model. On Hitchcock's view, taking 100 mg of the drug is a cause, if we are contrasting it with taking 0 mg, since causation should be understood as a ternary relation between C and E relative to some alternative cause C*. On Menzies' view, taking 100 mg is a cause if it makes a difference given the 'default causal model' for that situation.¹⁰ So if the patient would normally have been given the higher dose but, due to some medical error, received the lower dose, then taking the lower dose isn't a cause. But if the patient would not normally have received any dose of the drug at all, then receiving 100 mg of it is a cause. Either way, causal contextualism follows, as the truth-value of 'C causes E' is determined, in part, by contextual factors.

⁹ Here I follow Halpern and Pearl (2005), on whose view only endogenous variables can cause or be caused (p. 851). Both Hitchcock (2007) and Menzies (2007) present a slightly different account that does not have this implication.

¹⁰ Roughly speaking, the 'default causal model' is one in which all the values of the variables represent the 'normal or natural' states of the system, Menzies (2007, p. 208).

Causal contextualism, then, is well motivated. As well as being central to the plausible structural equations framework, it gains support from the fact that it offers a convincing diagnosis of a number of problematic cases. In what follows, I shall assume its truth in order to show how CC can aid us in an understanding of moral responsibility.

3 Resolving the Puzzles

We have, in CC, the resources to explain the intuitive data from §1 in a satisfying way.

3.1 Sharks and Missile

Recall that the difference in intuitive judgement concerning *Missile* and *Sharks* has moved some to suppose that responsibility for the consequences of acts and omissions must be treated asymmetrically. In particular, it has been used to motivate the view that responsibility for the consequences of omissions requires the ability to prevent the outcome, whereas responsibility for the consequences of acts does not.

Consider, first, *Missile*. We begin by considering whether Elizabeth is morally responsible for her act of launching the missile (according to the conditions of our preferred theory). Then we consider whether this act was a cause of the devastation of the city. Since this is a simple case of early pre-emption, such cases cause headaches for causal analyses, particularly those, like the structural equations approach, that appeal to the intuition that causes are what ‘make the difference’ to the effect. But all such sophisticated causal analyses have developed some manner of dealing with these cases to preserve the robust intuition that Elizabeth is the cause.¹¹ So granted that Elizabeth is morally responsible for launching the missile, CC is consistent with the intuition that Elizabeth is morally responsible for the devastation of the city.

What of *Sharks*? Again we can assume that John is morally responsible for his omission to act on behalf of the child. Now let’s suppose, to simplify matters, that given the presence of the sharks, the child’s death was a certainty, so no rescue attempt would have succeeded. In such a case, John’s failure to act does not raise the probability of the child’s death, it does not ‘make a difference’ to that outcome. We can model this by taking for granted the presence of the sharks by not having an endogenous variable corresponding to it in our model. Rather we include it, set at 0, in the set U of exogenous variables. As a result, whether or not John acts makes no

¹¹ See, for instance, Hitchcock’s (2001, §5) treatment of pre-emption cases, utilizing the structural equation approach and ‘ENF’ counterfactuals. The basic idea is that the route from Elizabeth to the bombing counts as ‘active’ (p. 286), since the bombing counterfactually depends upon Elizabeth’s actions, given that we hold the actual value of Carla’s actions fixed. So given that Elizabeth’s launching of the missile = 1, and Carla’s = 0, it must be true that had Elizabeth’s launching of the missile = 0, then the bombing would not have occurred.

difference to whether the child dies, and so his inaction is not a cause of the child's death.¹²

The point here is not dependent upon the particular details of some causal analysis or other, rather the thought is that any good analysis should preserve the intuition that John is not a cause of the child's death, whereas Elizabeth is a cause of the city's destruction. So CC gives us the right result without having to compromise the unity of our account of moral responsibility. In particular, there is no reason to suppose that responsibility for acts and their consequences, on the one hand, and omissions and their consequences, on the other, are asymmetric with respect to the ability to prevent the outcome. For this reason, CC has the advantage of simplicity over the asymmetric accounts proposed by Clarke and Sartorio.¹³

3.2 Two Buttons and Stuck Button

Two Buttons and *Stuck Button* raise significantly more complex issues than do *Sharks* and *Missile*. We begin, once again, by considering whether or not the agents in question were morally responsible for their failure to press the button. (Or, if we prefer, whether they were responsible for their decision not to press the button, of which their omission was a causal consequence.) So we might consider whether the agent's decision was the result of a moderately reasons responsive mechanism (see Fischer and Ravizza 1998), whether they identified with their decision (see Frankfurt 1971), and so on. Suppose that all such conditions are met and, moreover, that being clued up colleagues at the chemical factory, they are fully aware of the consequences of their decisions. Assume, further, that neither agent A nor B has any special knowledge regarding each other's character, which would lead them to suspect that their colleague was not an 'average' employee.¹⁴

The next question is whether their omissions should be considered causes of the explosion. This is when it matters what account of causation is offered. *Two Buttons* is a case of overdetermination and such cases create problems for any theory of causation based on the idea that causes make a difference to their effect. It's not true that, without any one cause, the effect would not have happened, but nevertheless, it looks like we have two causes of the same event, rather than no causation at all.

To preserve this intuition, proponents of the structural equations framework point out that, in overdetermination cases, the counterfactual dependence of the effect on any one cause is masked by its dependence on the other. So, by screening off the effect of one of the cause variables, we reveal the counterfactual dependence of the

¹² Alternatively, we may model the case differently, by using an endogenous variable for the sharks within the model (see footnote 21 for more details). What matters for our purposes is that *Sharks* is not modelled as a case of overdetermination (see pp. 15–16), since intuitively the presence of the sharks is seen as causally trumping John's inactivity.

¹³ In addition, Fischer and Ravizza present a complex analysis of moral responsibility for the consequences of actions (see 1998, p. 112), supporting it by showing how it can deal with four cases, *Missile* being one of them. But the complexity of their account can be dispensed with in favour of CC. The cases they cite all involve pre-emption or switching and can be dealt with given the approach above, but I will spare the reader the details. (For more on how to treat such cases, within the structural equations framework, see, for instance, Hitchcock 2001 §5 for pre-emption cases and §7 for switching cases.)

¹⁴ The relevance of this assumption will be made clear later in the discussion of *Lazy Two Buttons* (§4.3).

other, and vice versa.¹⁵ When we are dealing with singular causation, it is generally impermissible to fix the variables at non-actual values, since this will interfere with the causal relations. For example, suppose that in a set of equations modelling the fact that a match lit when struck, we set the variable for oxygen at its non-actual value, 0. This is clearly unacceptable, as it gets the causal structure of the actual situation wrong; the match wouldn't light. We avoid this outcome by permitting only non-actual values for those variables falling within, what Hitchcock and Woodward call, 'the redundancy range' (Hitchcock 2001, p. 290, Woodward 2003, p. 83). These are those variables which, if altered, will not change the value of the effect. So, unlike oxygen, in a case of overdetermination, since changing the value of one of the causes will not change the value of the effect, this variable falls within the redundancy range.

We can now apply this to the case of *Two Buttons*. Suppose that we use the following endogenous variables to model the situation:

BA = 1 if agent A's button is pressed, 0 if not.

BB = 1 if agent B's button is pressed, 0 if not.

E = 1 if explosion, 0 if not.

The set, U, of exogenous variables includes those conditions that are assumed to hold in that situation, e.g. that the buttons are working properly, that there is oxygen in the air, etc. Now we can model the situation as follows:

BA = 0

BB = 0

E = 0 iff BA & BB = 1

So this tells us that the explosion wouldn't have occurred if both A and B had pressed the button. It also tells us that the explosion would have occurred if, both A and B had not pressed the button, if A had pressed the button but B hadn't, and also if B hadn't pressed the button but A had. Given, then, that BB is set at 0, the value of BA makes no difference to E. However, since BB falls within the redundancy range for the effect, it is permissible to change it to its non-actual value, 1. Given this contingency, we see that E is counterfactually dependent upon BA, since if this were set at 1, the effect would not have occurred.

This, of course, does not solve our problem, since the same reasoning applies to the case of *Stuck Button*. Taking these as our variables:

BC = 1 if agent C's button is pressed, 0 if not.

BS = 1 if the button is working, 0 if not.

E = 1 if explosion, 0 if not.

BC = 0

BS = 0

E = 0 iff BC & BS = 1

¹⁵ See Halpern and Pearl (2005), Hitchcock (2001) and Woodward (2003, chapter 2) for the technical details and a formal definition of 'cause'. Here, I just present a very informal explication of the basic idea.

Since BS also falls within the redundancy range of the effect, it is permissible to change the value of this variable to 1. Given this contingency, we again see that the effect is counterfactually dependent upon the value of BC, since if this were set at 1, the effect would not have occurred. So if one is the cause, so is the other.

It is at this point, however, that the causal contextualism of the structural equations framework comes into play. Although the counterfactual structure of *Two Buttons* and *Stuck Button* is identical, questions concerning what causes what depend upon the right choice of model, and how we model the situation depends upon the context.

Consider, for instance, a classic problem case for omissions. Poppy promised to water my flowers while I was on holiday, but she failed to do it. Intuitively, her failing to water the flowers was a cause of their death since, if she had bothered to water them, they wouldn't have died. But it's equally true that if the Queen of England had watered my flowers, then they would have survived, and it seems wrong to view the Queen's failure as a cause of their death. The reason for this, according to Pearl and Halpern, is that the right model of the causal situation is one that includes an endogenous variable corresponding to Poppy's failure to water my flowers, but none corresponding to the Queen's failure to water my flowers.¹⁶ In this way, the Queen's non-activity is assumed as part of the background conditions. It is something that we reasonably take for granted and so, given the context, it cannot act as a difference maker for the effect. So we get the intuitive conclusion that only Poppy's failure to water my flowers, not the Queen's, is a cause of their dying.

How do we decide which is the right choice of model? Halpern and Hitchcock write,

Nature does not provide a uniquely correct set of variables. Nonetheless, there are a number of considerations that guide variable selection...inappropriate choices give rise to systems of equations that are inaccurate, misleading, or incomplete in their predictions of observations and interventions. (2010, p. 394)

Some possibilities do not count as 'serious possibilities,' for example the Queen's watering my flowers.¹⁷ But what counts as a serious possibility will vary depending upon the context. In a well-stocked, wealthy hospital where it is standard practice to issue 200 mg of the drug, the doctor's giving the patient 200 mg of the drug will count as a serious possibility. In a hospital in a destitute, war-torn country with no such medical supplies, it will not.

Various factors will determine which situations we take to be serious possibilities, and thus which deserve to be modelled. Key amongst them is what, in the situation, is 'normal' or to be expected. When we have a deviation from the

¹⁶ See Halpern and Pearl 2005, p.871. A similar solution is offered by Menzies (2007) which appeals to his aforementioned notion of default models. Roughly, given that Poppy (but not the Queen) promised to water the flowers, the default model of the situation is one that models Poppy's watering the flowers and them surviving. So when we plug in the actual values and note that the effect changes, we see that it was Poppy's not watering the flowers that made the difference in this situation.

¹⁷ See Woodward 2003 §2.8 for a discussion of 'serious possibilities'.

norm we look for some difference-maker to explain the deviation.¹⁸ How we understand what counts as the ‘norm’ will, again, vary depending on the context. Sometimes we will be concerned with statistical norms, i.e. what was statistically most likely in that situation. At other times, legal, social or moral norms will be paramount, as in the case of Poppy’s obligation to water my flowers.¹⁹ In other contexts, we will be concerned with the proper functioning of a system, for instance, the way the organ/machine should have worked. But, in addition, there may well be further factors that play a determining role. For instance, Woodward comments,

considerations having to do with whether an outcome is controllable at all (or easily or cheaply controllable) by current technology might also matter. It seems unlikely that there is any algorithm for determining whether a possibility will be or ought to be taken seriously. (2003, p. 88)

How do such considerations feed into an understanding of *Two Buttons* and *Stuck Button*? In *Two Buttons* it is reasonable to assume that it is possible to manipulate both agents, A and B, through the usual array of incentives and punishments. As they are both factory employees whose job description includes pressing buttons in an event of an emergency, it is normal (socially, morally, legally and presumably statistically) for them to do so. Consequently, their both doing so is regarded as a serious possibility. As there is nothing to choose between agents A and B, their failures to press their buttons will each be values of an endogenous variable in the appropriate model. Hence, we may follow the earlier reasoning and get the hoped for result that they are both causes of the explosion.²⁰

¹⁸ Following Hart and Honoré (1985), this aspect has been recently emphasized by Menzies (2007, 2009), Hall (2007) and Halpern and Hitchcock (2014). It is also supported by studies of people’s actual judgments concerning causation. Hitchcock and Knobe (2009) report on a series of experiments which purport to show that we are much more likely to judge an event as the cause if it violates some norm.

¹⁹ On this way of viewing causation, there is no reason to deny that, sometimes, ethical factors (e.g. whether or not the person had an obligation to perform a certain action) determine the right choice of model in a particular situation. This, however, doesn’t render causation thoroughly subjective for two reasons. First, the counterfactual structure of the situation in question that, in large part, determines the causal facts is free from such normative elements. So, on this view, there are various objective facts regarding difference making relations, and these can be mapped in different ways, depending upon the context, to give us our causal claims. What causation is tracking, then, is an independent relation out there in the world (hence the claim that causation is not ‘thoroughly subjective’), even though, when we make causal claims, we do have to model that information in some way, and so our causal claims do reflect our interests.

Second, in many areas (e.g. physics), any ethical considerations would be out of place (although this is not to deny that subject-specific norms may well be relevant to the choice of model). In this way, we can hold onto the idea that causation is necessary for moral responsibility for consequences, whilst acknowledging that ethical or broader normative considerations do, sometimes, feature in the correct model determining what we are causally responsible for.

²⁰ On Menzies’ view (2007), the default model for the situation is one where both variables BA and BB are set at 1, and so the explosion does not occur. So when we come to modelling the actual situation, given the added complication to deal with overdetermination cases, BA and BB are both what make the difference to the effect.

Contrast this with *Stuck Button*. Here, the automated system is not like agent B. There are relevant differences which, I suggest, make it at least reasonable for the modeller to use a different set of equations. One difference between these cases is that, as we assumed, the agents in *Two Buttons* could easily have pressed the button. Either may have changed their mind at the last minute. By contrast, system deficiencies or violations of proper functioning appear more entrenched. It seems much less likely, without further intervention from a mechanic, that the stuck button would spontaneously have become unstuck. The default state of the automated safety system, given that it is broken at time t , is for it to remain so until it has been fixed. As a result, we are much less likely to regard the spontaneous unsticking of a jammed button as a serious possibility than we are the spontaneous decision of agent B to press the button.

This difference arguably makes appropriate different models for the situations. So, in *Stuck Button*, we take for granted the malfunctioning of the button by not having an endogenous variable corresponding to it in our model. Rather we include it, set at 0, in the set U of exogenous variables. As a result, whether or not agent C presses the button makes no difference to whether the explosion occurs, and so is not a cause of the explosion. In *Two Buttons*, on the other hand, we have granted that both A's and B's pressing the button are considered serious possibilities. So, as I said above, this is represented in the model by two endogenous variables.²¹

One might object that the analysis does not present us with any justification of our intuitions, just ad hoc rationalisations of them. Worse than this, there is no reason to endorse the causal model that explains the intuition, other than the fact that we are attempting to preserve the intuition. So there's an inherent vicious circularity at the heart of these rationalisations—we can explain our intuitions only by assuming them to be true.²²

This objection can, I think, be dealt with at least given a certain understanding of what the project is. The proposal is trying to explain some data about moral

²¹ On Menzies' view, the difference between the two cases is arguably not marked, as it seems intuitive to say that the default model is one where both BS and BC are set at 1. We could however argue, for the reasons given above, that in this situation the default model should assume only the proper functioning of the system, not the action of the agent. In this, it is like McDermott's example of a fielder catching a cricket ball. Between the fielder and the window stands a high wall (1995, p. 525). Does that fielder prevent the ball from hitting the window? Intuitively not, as the wall would have stopped the ball from hitting the window anyway. To capture this situation, on Menzies' view (see 2007, §5B), we contrast the actual values with those of the default model, which include the presence of the wall and the absence of the fielder's catch. Since the fielder's catch makes no difference to the window's breaking, it is not a cause. We might argue that this is the right way to model the situation, even though it was the fielder's function, as a player in the game, to catch the ball, since the wall is such an entrenched and effective preventer in that situation. Such considerations can similarly motivate a different default model in *Stuck Button* (and *Shark*), the idea being that the agent's omission in this case is like the fielder (John's inaction), whose potential causal input has been trumped by the presence of the stuck button (shark), which is analogous to the wall. This might well be a preferable way of modelling the situation, to that suggested above, since it would allow us to say that, relative to this model, the stuck button (shark) is the cause of the explosion (child's death) since this is what makes the difference with respect to the default model. On Pearl and Halpern's approach, in contrast, we would need to remodel the situation so that the Stuck Button (shark) was an endogenous variable in the model to get this result (see Halpern and Pearl 2005, Example 5.6 for more details).

²² Many thanks to an anonymous referee for this objection.

responsibility which it assumes is intuitive. In this sense, it is engaging in reflective equilibrium: it is attempting to find a theoretical underpinning for judgements that we already, but fallibly, take to be true. So I am assuming that certain judgments regarding our responsibility for consequences are, defeasibly, true and seeking a way of accommodating and explaining them within a theory of moral responsibility. But, and this is key to the reply, I needn't assume that, for instance, agent A is morally responsible for the explosion in order to justify the choice of causal model. As I see it, in that scenario, there are facts about difference making relations, statistical probabilities, social and moral norms, the abilities of agent A etc. which can all be used to justify the claim that there is an appropriate causal model in which agent A is a cause of the explosion independent of any assumptions regarding her moral responsibility for the explosion.

But where, we might persist, is the *justification* of the fact that agent A is morally responsible for the explosion? On this view, the justification comes at the level of the choice of causal model. What the correct causal model is, and to what extent the context is able to fix upon any one model, will often be a matter of dispute. But what is key, as Halpern and Hitchcock comment, is that we are

able to justify the modelling choices made. A lawyer in court trying to argue that faulty brakes were the cause of the accident needs to be able to justify his model; similarly, his opponent will need to understand what counts as a legitimate attack on the model. (2010, pp. 384–385)

What it takes to defend our modelling choices might well depend upon domain specific norms and expectations, as well as details concerning what it is reasonable to judge will be the case, given those specific circumstances. But the thought is that by defending our modelling choice, we will thereby get our justification of the fact that agent A is morally responsible for those consequences. Moreover, there is no reason to think that this justification will involve appealing to the very fact that agent A is morally responsible for the explosion. Rather it will appeal to the aforementioned facts about difference making relations, details about the specific circumstances, the social and moral norms, the agent's abilities, etc. This, once combined with CC and a theory of moral responsibility for acts and omissions, gives us the intuitive result we are after.

3.3 Lazy Two Buttons and Chancy Stuck Button

It may be objected that this conclusion cannot be robustly defended, since it is subject to how we model the situations and on what counts as a 'serious possibility'. Consider the following variations on the two cases:

Lazy Two Buttons: A dangerous chemical is leaking which, if unchecked, will lead to a massive explosion. Agent A and agent B can seal off the leak by simultaneously pressing their emergency buttons. Both decide to do nothing and an explosion ensues. In addition, B is the most immoral, lazy agent imaginable and there is no chance that she would have pressed her button.

Chancy Stuck Button: A dangerous chemical is leaking which, if unchecked, will lead to a massive explosion. Two buttons need to be pressed simultaneously to seal off the leak. Agent C must press one emergency button; the other button should automatically be activated once a leak has been detected. Unfortunately, the button on the automated system is jammed, although there was a 30% chance that it would have worked properly. In addition, agent C decides not to press her button.

It seems that in *Chancy Stuck Button*, there is a serious possibility that the button would have been pressed. In *Lazy Two Buttons*, on the other hand, B's pressing the button is not a serious possibility. With these cases we will be unable to preserve our initial intuitive judgements that both A and B are responsible for the explosion whilst C is not.

This objection is not, however, compelling. In *Lazy Two Buttons*, when we construct the appropriate causal model, the fact that agent B should have (legally, morally, as part of her job) pressed the button outweighs the statistical unlikelihood of her actually doing so. So agent A and agent B are still appropriately considered causes of the effect, so morally responsible for the explosion. Interestingly, this is not so if agent A knows B's character, and so knows that there is no point in her pressing the button.²³ We can explain this shift by saying that, given the context, agent A models the situation so that her inaction makes no difference to the effect, since the inaction of agent B is taken as a given. If agent A can, in Halpern and Pearl's words, 'defend her choice of model' (2005, p. 871), she would not be morally responsible for the explosion, since she wouldn't count as a cause of the explosion.

What of *Chancy Stuck Button*? Here I think our intuitions regarding agent C's moral responsibility for the explosion are influenced by such changes. As with *Full Shark*, one might argue that agent C is morally responsible for the explosion. Since there was some chance that the other button would have been depressed by the automated system, her inaction did raise the chance of the explosion, and so made a difference to the likelihood of the effect. Consequently, the button's malfunctioning should not be represented as a given, exogenous causal variable.

Depending upon how the details are further spelt out, however, our intuitive reactions may well be more nuanced than this suggests. If there is any time, *t*, before the explosion, at which the button is definitely jammed, then we might say that agent C's pressing the button made no difference to the effect since it did not increase its probability.²⁴ The actual malfunctioning of the system at *t*, even though not guaranteed at a time before *t*, trumps the inactivity of the agent. However, if there remains a 30% chance that the second button will be activated right up until the time of the explosion, then agent C's inactivity does increase the likelihood of the effect.

²³ Unlike in the case of a technical malfunction, we may question whether any such belief about a person's character and their future actions could justifiably be considered certain enough. But it is important to note that there may be contexts where the statistical unlikelihood of an agent's doing something outweighs factors concerning what legally, morally etc. they should do when formulating the appropriate causal model (see §3.4 for further discussion).

²⁴ See Hitchcock (2004) for a defence of this line within the structural equations approach.

3.4 Penned-in Sharks

This diagnosis of the difference between *Two Buttons* and *Stuck Button* compares favourably to those presented by either Sartorio (2005) or Byrd (2007), mentioned in §1, since it offers a principled analysis of why agents A and B are morally responsible whilst C is not.²⁵ The difference arises from the simple fact that, in *Two Buttons*, A is a cause of the explosion, and so morally responsible for this outcome. In contrast, this is not the case in *Stuck Button*. Moreover, these causal judgments can be justified given a well-supported approach in the philosophy of causation.

CC also has the further advantage of mirroring ambivalent reactions to certain tricky cases. Consider this case,

Penned-in-Sharks: The evil Emery will release a deadly school of sharks if anyone attempts to save the drowning child. But John decides not to jump in, completely unaware of the dangers of doing so. (Based on Fischer and Ravizza 1998, p. 138)

Fischer and Ravizza argue that, unlike in *Sharks*, John is morally responsible for the child's death in *Penned-in-Sharks*. This judgment has been controversial.²⁶ It is difficult to see any reason for distinguishing between the two cases, at least in their present bare form—difficult, but not impossible. Given that, as in *Two Buttons*, another agent's decision is involved in bringing about the terrible consequences, we may be inclined to judge that factor as susceptible to spontaneous change. It seems a less entrenched feature of the situation than the man-eating quality of the circling sharks. We can model this by making Emery's releasing of the sharks an endogenous variable:

RS = 1 if sharks released, 0 if not.

JA = 1 if John acts, 0 if not.

CD = 1 if child dies, 0 if not.

RS = 0

JA = 0

CD = 0 iff RS = 0 & JA = 1

Since the sharks being released would be subsequent to any lifesaving attempts on the part of John or others, when John makes his decision not to act, the releasing of the sharks is set to its actual value, 0. As this is held fixed when assessing whether John's inaction made a difference, we get the result that John's inaction is a cause of the child's death, since if JA were set at 1, the child would not have died.

Although it is unclear whether this is the most natural way of modelling the situation, further specifications of the case might make it so. If we slightly alter the example, so that Evil Emery would only *probably* release the sharks, then it seems much more reasonable to view John's omission as a cause of the child's death. This is what occurs when we switch from *Sharks* to *Full Sharks*. The same principled

²⁵ Furthermore, as I also pointed out in §1, neither Clarke's (1994) nor Sartorio's (2005) accounts can deliver the intuitively correct result for *Two Buttons*.

²⁶ Byrd, for instance, agrees with Byrd (2007, p. 63), whilst Clarke disagrees (2011, pp. 609–610).

explanation can be given of both changes in assessment: since the action of Emery/ Sharon the shark is less certain, their inaction becomes a much more serious possibility, thus justifying a differing causal model to that of *Sharks*.

But what if it is guaranteed that Emery will release the sharks? Then it does seem counter-intuitive to say that John is morally responsible for the child's death in *Penned-in-Sharks*, but not in *Sharks*. By using a different set of endogenous variables to model the causal situation, this intuition can be respected.

EP = 1 if Emery plans to release the sharks, 0 if not.

JA = 1 if John acts, 0 if not.

CD = 1 if child dies, 0 if not.

EP = 1

JA = 0

CD = 0 iff EP = 0 & JA = 1

Given a set of exogenous variables, U, which guarantee the success of Emery's plan, since the value of EP = 1 at the time when John does nothing, the value of JA makes no difference to the child's death. Consequently, we can say that John is not morally responsible for the child's death in *Penned-in-Sharks*, because he is not a cause of that death.

It might be objected that this is inconsistent with the judgment given in *Lazy Two Buttons*. There I claimed that since agent B should have (legally, morally, as part of her job) pressed the button, this outweighs the statistical unlikelihood of her actually doing so. But the same reasoning applies equally here. We can stipulate that Emery (legally, morally, as part of his job) shouldn't release the sharks. So the causal model should reflect this, rendering John once again a cause.

This objection fails to recognise that *Lazy Two Buttons* and *Penned-in-Sharks* are analogous in just the right ways. Suppose that agent A of *Lazy Two Buttons* knows that there is absolutely no chance (imagine an omniscient God tells her) that agent B will press the button. Then she can sit back, despairing—her inaction isn't a cause of the explosion, there is absolutely nothing that she can do. Similarly suppose that John knows, with absolute certainty (God again), that Emery will release the Sharks. (And knows, again with equal certainty, that the sharks will gobble him up before he can save the child.) Then he would be a fool or suicidal to jump in, as two lives would be lost rather than one. His inaction isn't a cause, since there is nothing he can do which will make a difference to that dreadful outcome. In both cases, it is up to the modeler to 'defend her choice of model' (Halpern and Pearl 2005, p. 871) and we should model the situations in different ways, with differing outcomes, depending upon the details of the cases.

The current proposal then, has the significant advantage of accounting for differing reactions to various cases of moral responsibility for omissions.²⁷ Moreover, it can explain shifting intuitions to these cases, given various ways of further specifying them.

²⁷ For another example of divided opinion, see McIntyre's case (1994, p. 485). Such cases can be treated in a similar way.

4 An Objection

Despite these advantages, one might argue that the proposal has an Achilles' heel. If causation is contextual, and moral responsibility is based on these causal claims, then ascriptions of moral responsibility for consequences are also only contextually true or false. This implication is unacceptable. Take, for example, *Sharks*. There is a possible way of modelling that situation, the one outlined, which results in John's not being morally responsible for the child's death. But there are other ways we might model the situation that do not have this result. We might, for instance, instead hold fixed the fact that the sharks didn't move towards the child at t_1 , in which case, the sharks wouldn't have prevented the rescue, and John would have been morally responsible for the child's death.

I accept that, on this analysis, ascriptions of moral responsibility for consequences are contextually true or false, but I deny that this is unacceptable. The objection assumes that such contextualism results in rampant subjectivity—that anything goes. But this misconstrues the nature of the causal contextualist's position. In many, if not all, cases, there will be the best or most appropriate way of modelling the causal situation.²⁸ To illustrate, consider again the case mentioned earlier, 'Flora's failure to water my flowers caused them to die'. Having an endogenous variable corresponding to the Queen's failure to water my flowers would be an inappropriate causal model. So in that context (and all relevantly similar contexts) it is false to say that the Queen was a cause of my flower's dying. Similarly, in *Sharks*, there would need to be some reason to hold the movements of the sharks fixed at t_1 for this to be the appropriate causal model. Given that sharks do not generally stay motionless in the water (and there is nothing guaranteeing this in the story, e.g. they have not been disabled by a stun gun or drugged), this does not get the causal model right.

This disarms the objection. Since the appropriate causal model is constrained by the circumstances, so is what we are morally responsible for. Furthermore, there is good reason to think that we should be assessing moral responsibility for consequences within contexts, since if we change contexts, facts that are relevant to what an agent is morally responsible for also potentially alter. For example, in *Full Shark*, it is highly unlikely that the child would have been attacked, so the causal model of the situation should be different from that of *Sharks*. But our intuitions regarding whether it is true to claim that, 'John's inaction caused the child's death' and 'John is morally responsible for that child's death' also shifts in these two cases.

One might worry that, since the context is not under an agent's control, what the agent is morally responsible for is subject to an unacceptable degree of moral luck. But any analysis of moral responsibility concerned with the consequences of our acts and omissions has to allow that what they cause is affected by factors outside our control. (E.g. cf. Zimmerman's 2002, example, of a passing bird taking the force

²⁸ The contextual facts should not be taken to include only what nature provides (which, given the kinds of cases under discussion, are highly unlikely to result in one most appropriate causal model), but also facts about e.g. moral, social, legal etc. norms. For more discussion of this point, see Woodward (2003 §2.8) and Halpern and Hitchcock (2010, 2014).

of a would-be assassin's bullet rather than the planned human target.) The only way of shielding us from this resultant moral luck is by limiting ascriptions of moral responsibility to our acts and omissions, dumping causal contextualism will not help.²⁹ But perhaps, it might be argued, this is good reason to adopt a more radical view? If the only difference between agent A's and C's blameworthiness rests on some technical malfunction, surely this is deeply unfair?³⁰

I think we should proceed with caution here, however. To say that agent A, unlike agent C, is responsible for an explosion and the resulting deaths, does not commit one to any claim about the severity of punishments agents A and C should receive. It doesn't even commit one to the view that agent A is more blameworthy than agent C, since from the claim that agent A is morally responsible for more bad events than agent C, it does not follow that agent A is more blameworthy than agent C. So whilst I do think that we are responsible for the consequences of our acts and omissions, and this is subject to resultant luck, this claim is only subject to the objection that it is unfair or unjust if we make additional assumptions regarding how the agent should be treated.³¹

Still, one might object, according to causal contextualism there may be no most appropriate causal model, given the facts of the situation. Rather a number of different causal models may fit the context equally well. If this were so, then an agent may be morally responsible for a particular consequence, given certain ways of modelling the situation, but not on others. The resulting indeterminacy would not be merely epistemic in character. It wouldn't simply be that we do not *know* whether the agent's act or omission caused the later events, but rather that there is no fact of the matter concerning whether or not they caused it, and so whether or not they were responsible for it. This metaphysical indeterminacy, it might be thought, is worrisome.

To develop this objection, we would need to know why this metaphysical indeterminacy is supposedly problematic. If we are focusing on the practical implications of the view, the consequence of this metaphysical indeterminacy is no different from that of epistemic indeterminacy. If we do not know what the agent's act or omission caused, then to avoid any unfairness, we should limit our moral appraisal to what we do know, namely, to the nature of their act or omission. The same holds if there is no fact of the matter concerning what they caused. Given mere epistemological indeterminacy, an omniscient God would know precisely what we were all morally responsible for. But such supernatural considerations fall outside of

²⁹ Even this much is unclear since it depends upon your theory of action. On many theories, what act we perform is determined partly by its consequences (hence the act was a killing, rather than a mere attempted killing, because of what happened). In this way, it might be that CC becomes involved in an understanding of what act has taken place.

³⁰ Many thanks to an anonymous referee for this comment. See Graham (2014) for a defence of an alternative way of sorting out the problem—making all the puzzles dissolve by denying that we are morally responsible for the consequences of our acts and omissions.

³¹ It is clear that our current legal system does treat people differently depending upon whether or not they successfully completed their crime, and I agree that there are reasons to worry about this (see Feinberg 2003 for an interesting discussion). But this is not a commitment of the view I am defending.

the practical implications of the view for our own practices of holding people to account.

It might be suggested that although, practically speaking, the import of the view is not problematic, metaphysically it is. A discussion of this point lies outside the bounds of this paper. As I made clear earlier, I am assuming the truth of causal contextualism, and the possibility of such metaphysical indeterminacy is a commitment of this view. Since it is highly plausible to say that moral responsibility for consequences has something to do with causation, if causal contextualism is true (and it is a well-supported theory) we are going to have to bite this bullet anyway.³²

But, having noted this defensive strategy, I think that the best response to the worry is the admittedly bold one of taking such metaphysical indeterminacy to be an advantage of the view. Although it is an open empirical question whether there are any actual cases where the facts fail to secure one most appropriate causal model for the situation, there seem to be conceivable cases of which we might want to say that, when all the evidence is in, it is indeterminate what the agent is morally responsible for.³³ Indeed, one might argue that this is true of John in *Penned-in Sharks*, given certain specifications of the example. In such cases, the analysis has the advantage of both predicting that our intuitions will be ambivalent, and providing an explanation of this fact: they are unstable because the facts fail to secure a best causal model of the situation.

5 Conclusion

The puzzles that arise when attempting to deal with the consequences of acts and omissions do not stem from anything peculiar about moral responsibility. Rather, they are inherited from the subtlety of our causal judgements. By tying our theory of moral responsibility to causal contextualism, we simplify our account of responsibility.

It might be objected that since causal contextualism fails to give us a fixed, determinate set of objective causal facts, unlike non-contextualist theories of causation, appealing to it ultimately results in a more complex account.³⁴ This objection raises difficult questions regarding which accounts are more complex, unified, etc. But I have been assuming throughout that we need causal contextualism to account for all of the tricky causal data. Consequently, we might as well use this to our advantage elsewhere. So my claim is not that CC plus causal contextualism is

³² Note: we needn't spell out the relationship between causation and moral responsibility for consequences in terms of CC for this to follow. See Sartorio 2007 for further ways of construing that relationship—which she labels 'entailment dependence' (p. 750), 'quasi-entailment dependence' (p. 753), 'entailment-reverse dependence' (p. 755), and 'transmission dependence' (p. 761)—all of which would have this implication.

³³ An anonymous referee kindly notes that this view may have an ally in Aristotle, who thinks, roughly, that knowledge of the Good does not admit of a high degree of precision or accuracy. Aristotle writes, 'it is a mark of an educated person to look in each area for only that degree of accuracy that the nature of the subject permits. Accepting from a mathematician claims that are mere probabilities seems rather like demanding logical proofs from a rhetorician.' (Nicomachean Ethics, book 1, 1094b).

³⁴ Many thanks to an anonymous referee for this objection.

simpler than CC (or some other causal condition, see footnote 32), plus an asymmetry thesis, plus non-causal contextualism. Rather it is that CC plus causal contextualism is simpler than some causal condition, plus an asymmetry thesis, plus causal contextualism.

And, finally, why is CC so compelling? Because it provides a clear, principled explanation of why the agent's moral responsibility extends to just those events that it does. I am morally responsible for an outcome if I am morally responsible for the act or omission that caused it. Rival accounts struggle to explain the extent of our moral responsibility. The fact that CC, alongside causal contextualism, allows us to accommodate the intuitive data in a simple and unified way strongly recommends it as a condition on responsibility.

Acknowledgements Many thanks to the anonymous referees of this journal, who gave many helpful comments, and the participants of a graduate workshop in Manchester. Also thanks, as always, to Joel Smith. Work on this paper was generously funded by a grant from the Arts and Humanities Research Council.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Aristotle, translated and edited by R. Crisp (2000) *Nicomachean Ethics*, Cambridge University Press.
- Byrd, J. (2007). Moral responsibility and omissions. *Philosophical Quarterly*, 57, 56–67.
- Clarke, R. (1994). Ability and responsibility for omissions. *Philosophical Studies*, 73, 195–208.
- Clarke, R. (2011). Omissions, responsibility, and symmetry. *Philosophy and Phenomenological Research*, 82, 594–624.
- Clarke, R. (2014). *Omissions: agency, metaphysics and responsibility*. Oxford: Oxford University Press.
- Cooper, D. E. (1968). Collective responsibility. *Philosophy*, 43, 258–268.
- Feinberg, J. (2003). 'Criminal Attempts' in his *Problems at the Roots of Law* (pp. 77–102). Oxford: Oxford University Press.
- Fischer, J. M., & Ravizza, M. (1998). *Responsibility and control: A theory of moral responsibility*. Cambridge: Cambridge University Press.
- Frankfurt, H. (1971). Freedom of the will and the concept of a person. *Journal of Philosophy*, 68, 5–20.
- Frankfurt, H. (1994). An alleged asymmetry between actions and omissions. *Ethics*, 104, 620–623.
- Graham, P. (2014). A sketch of a theory of blameworthiness. *Philosophy and Phenomenological Research*, 88, 388–409.
- Hall, N. (2007). Structural equations and causation. *Philosophical Studies*, 132, 109–136.
- Halpern, J., & Hitchcock, C. (2010). Actual causation and the art of modelling. In R. Dechter, H. Geffner, & J. Halpern (Eds.), *Heuristics, probability, and causality: A tribute to Judea Pearl* (pp. 383–406). London: College Publications.
- Halpern, J. Y., & Hitchcock, C. (2014). Graded causation and defaults. *British Journal for the Philosophy of Science* (advanced access).
- Halpern, J. Y., & Pearl, J. (2005). Causes and explanations: A structural-model approach. Part I: Causes. *British Journal for the Philosophy of Science*, 56, 843–887.
- Hart, H. L. A., & Honoré, T. (1985). *Causation in the law* (2nd ed.). Oxford: Oxford University Press.
- Hitchcock, C. (1996). The role of contrast in causal and explanatory claims. *Synthese*, 107, 395–419.

- Hitchcock, C. (2001). The intransitivity of causation revealed in equations and graphs. *Journal of Philosophy*, 98, 273–299.
- Hitchcock, C. (2004). Routes, Processes and Chance-lowering Causes. In P. Dowe & P. Noordhof (Eds.), *Cause and Chance: Causation in and Indeterministic World* (pp. 138–151). London: Routledge.
- Hitchcock, C. (2007). Prevention, preemption, and the principle of sufficient reason. *Philosophical Review*, 116, 495–532.
- Hitchcock, C., & Knobe, J. (2009). Cause and norm. *Journal of Philosophy*, 106, 587–612.
- McDermott, M. (1995). Redundant causation. *British Journal for the Philosophy of Science*, 46, 523–544.
- McIntyre, A. (1994). Compatibilists could have done otherwise: Responsibility and negative agency. *Philosophical Review*, 103, 453–488.
- Menzies, P. (2004). Difference-making in context. In J. Collins, N. Hall, & L. Paul (Eds.), *Causation and counterfactuals* (pp. 139–180). Cambridge: MIT Press.
- Menzies, P. (2007). Causation in context. In H. Price & R. Corry (Eds.), *Causation, physics, and the constitution of reality: Russell's republic revisited* (pp. 191–223). Oxford: Oxford University Press.
- Menzies, P. (2009). Platitudes and counterexamples. In H. Beebe, P. Menzies, & C. Hitchcock (Eds.), *The Oxford handbook of causation* (pp. 341–367). Oxford: Oxford University Press.
- Pearl, J. (2000). *Causality: Models, reasoning and inference*. Cambridge: Cambridge University Press.
- Sartorio, C. (2004). How to be responsible for something without causing it. *Philosophical Perspectives*, 18, 315–336.
- Sartorio, C. (2005). A new asymmetry between actions and omissions. *Noûs*, 39, 460–482.
- Sartorio, C. (2006). Disjunctive causes. *Journal of Philosophy*, 103, 521–538.
- Sartorio, C. (2007). Causation and responsibility. *Philosophy Compass*, 2, 749–765.
- Schaffer, J. (2005). Contrastive causation. *Philosophical Review*, 114, 327–358.
- Schaffer, J. (2012). Causal contextualisms. In M. Blaauw (Ed.), *Contrastivism in philosophy: New perspectives* (pp. 35–63). London: Routledge.
- Whittle, A. (2016). Ceteris Paribus, I could have done otherwise. *Philosophy and Phenomenological Research*, 92, 73–85.
- Woodward, J. (2003). *Making things happen*. Oxford: Oxford University Press.
- Zimmerman, M. (2002). Taking luck seriously. *Journal of Philosophy*, 99, 553–576.