

Manipulation and Independence

D. Justin Coates

1. Wisdom from the Past

Anticipating an important but otherwise recent turn in the arguments for the incompatibility of moral responsibility and causal determinism, John Wisdom (1934) invites us to consider the following question.

Suppose that all your acts are determined by your decisions, and your decisions by your knowledge (no doubt imperfect) of the consequences of your acts together with your desires for these consequences. But suppose that on the occasion of each decision the strength of your various desires is fixed by the Devil. Suppose that you float a bogus company and ruin thousands. Are you to blame? [Wisdom 1934, 116].

In response to this question Wisdom expresses skepticism.¹ Though it is true, he claims, that in such a case your act was “due to your nature,” *and* you had “freedom of self-determination,” *and* “your act was due to that part of your nature which is called your will,” *and* you had “freedom of self-direction” (Wisdom 1934, 117), it is dubious that you would be *ultimately* responsible for it or *really* to blame. More is required for moral responsibility than can be had given the truth of demonic intervention or, Wisdom suggests, causal determinism. Accordingly, Wisdom concludes that

You can, then, be to blame for an act only in so far as that act is caused by your nature, and, in particular, by your will; further, your will must be at least part of the *ultimate* cause of your act. To say that your will is part of the ultimate cause of your act is to say that while your will determined your



Journal of Information Ethics / Volume 28, Number 1 / Spring 2019 / pp. 115–130 /
ISSN 1061-9321 / eISBN 978-1-4766-2189-0 / © 2019 McFarland & Company, Inc.

act that will is not in its turn completely determined by something which is not your will [Wisdom 1934, 118].

Wisdom's view, then, is that to be morally responsible for your action—to be *to blame* in his parlance—an agent must be the *source* of her action in a metaphysically robust sense that is precluded by causal determination.

2. Returning to the Source

This thought lay dormant for some time in the philosophical debates concerning whether moral responsibility was compatible with causal determinism. For a long time the battle between compatibilists and incompatibilists raged over whether *free will*, which has traditionally been understood as the ability to do otherwise, was compatible with determinism. Incompatibilists denied that it was, and so concluded that moral responsibility—widely regarded as requiring this ability—was not possible in a deterministic system.² Compatibilists, on the other hand, offered analyses of the ability to do otherwise that were compatible with determinism and so concluded that the truth of determinism posed no threat to responsible agency.³

However, the battleground between compatibilists and incompatibilists has largely shifted in ways that make Wisdom's argument particularly prescient. An early harbinger of this shift was Harry Frankfurt's classic paper "Alternate Possibilities and Moral Responsibility" (1969), which purported to show that the ability to do otherwise was not necessary for responsibility.⁴ Though the true "moral" of Frankfurt's story is still a matter of significant controversy, this paper does give us reason to think that the emphasis placed on the ability to do otherwise is if not wholly misplaced, certainly overblown. Yet even this result will not secure victory for the compatibilist. For even if Frankfurt was right to think that moral responsibility does not entail the ability to do otherwise (however one analyzes that concept), there are, as Wisdom knew, good reasons to adopt incompatibilism.

John Martin Fischer (1982) makes this point by arguing (in essence) that incompatibilists chastened by the lessons of Frankfurt can return to the insight found in Wisdom.⁵ More precisely, Fischer claims that determinism might threaten moral responsibility not because it rules out otherwise open alternatives but because it, like the Devil in Wisdom's example, interferes with the *actual causal sequence* that issues in agents' actions. In other words, determinism might threaten moral responsibility because it precludes agents (or their wills) from being the causal sources of their actions in the way that responsibility requires. In the wake of this thought the kind of *source incompatibilism* that was suggested in Wisdom's work on free will and moral responsibility reemerged onto the scene. And these days, it has become the biggest game in

town, buttressed by arguments that are sophisticated and well-crafted descendants of Wisdom's comparatively modest musings on what the Devil's presence in the causal history of our action might mean for our theory of moral responsibility.

3. Against Compatibilism

Contemporary defenses of source incompatibilism are varied, but most of them rely on scenarios meant to elicit the very judgments that Wisdom found so salutary when he reflected on the Devil's role in your decision to float a bogus company and in the process, ruin thousands. From these judgments more careful arguments are constructed. Arguments of this general type are usually called *manipulation arguments*, though within that broad category there is wide diversity in the details of how these arguments are run.

3.1. *The Zygote Argument*

One important version of this kind of argument is Al Mele's (1995, 2006) *original design argument*, which Mele calls the "Zygote Argument." The Zygote Argument too has variants, but each starts with roughly the following scenario:

[A goddess named] Diana creates a zygote *Z* in Mary. She combines *Z*'s atoms as she does because she wants a certain event *E* to occur thirty years later. From her knowledge of the state of the universe just prior to her creating *Z* and the laws of nature of her deterministic universe, she deduces that a zygote with precisely *Z*'s constitution located in Mary will develop into an ideally self-controlled agent who, in thirty years, will judge, on the basis of rational deliberation, that it is best to *A* and will *A* on the basis of that judgment, thereby bringing about *E* [Mele 2006, 188].

Here we have a zygote that develops into a human agent named Ernie who is not that different from any of the rest of us in terms of his possession of the rational powers and capacities that compatibilists take to be the grounds of moral responsibility. Yet from this scenario, Mele constructs an explicit argument against the idea that the powers and capacities compatibilists regularly regard as the grounds of moral responsibility are not up to the task. It runs as follows.

1. Because of the way his zygote was produced in his deterministic universe, Ernie is not a free agent and is not morally responsible for anything.
2. Concerning free action and moral responsibility of the beings into whom the zygotes develop, there is no significant difference between the way Ernie's zygote comes to exist and the way any normal human zygote comes to exist in a deterministic universe.

but reasons-responsive process of deliberation, and he has the specified first and second-order desires. The neural realization of Plum's reasoning process and decision is exactly as it is in Cases 1–3; he has the general ability to grasp, apply, and regulate his actions by moral reasons, and it is not because of an irresistible desire that he decides to kill [Pereboom 2014, 79].

As I understand them, Cases 2–4 appear to demonstrate rather than simply assert the thought that with respect to moral responsibility, there are no important differences between manipulation of certain sorts and causal determinism. After all, if the judgment that Plum is not responsible in Case 1 is reasonable, then it is hard to see why one should not come to the same conclusion in Case 2. Mere temporal distance of the sort introduced in Case 2 between the neuroscientists' programming and Plum's actions certainly cannot be relevant for evaluating his responsibility. This means that the *explanation* for why Plum is not responsible in Case 1 must also coincide with the explanation for why Plum is not responsible in Case 2. It must be that in Case 2, it is still the case that Plum is not the source of his actions in a metaphysically robust sense. And if that is right, then it does not bode well for the compatibilist when we turn to less "fantastic" cases—cases like Cases 3 and 4 that might describe the way things are in our own world.

Reflection on Case 3, taken in isolation from Cases 1 and 2, does not so clearly result in the judgment that Plum is not responsible. It is true that he is determined by his upbringing, but it is possible that we too are determined by the practices and values that we inherited early on in our lives. And this does not obviously obviate our status as responsible agents. Nor does the bare intuition of non-responsibility on the part of the incompatibilist in this case (again, taken in isolation of Cases 1 and 2) amount to an argument (or a premise in an argument) for incompatibilism, since this is precisely the judgment that compatibilists (and maybe even some agnostics) would reject.

However, this is not the reason Pereboom introduces Case 3. Even if you initially lack the intuition that Plum is responsible in Case 3 and so you are open to compatibilism, you will not be "in the clear" dialectically speaking. The task for the compatibilist who wants to insist that Plum would be responsible in more ordinary scenarios like Case 3 is to "adduce a feature of [Plum-in-Case 3's] circumstances that would explain why he is morally responsible here but not in Case 2," (Pereboom 2014, 78). But it is not clear what that feature might be, since although the causal history of how Plum comes to decide to kill White is different in Cases 2 and 3, it nevertheless comes to the *same thing*, for the *same reason*, chosen by *otherwise identical counterparts*. Indeed, because the very features that explain why it is plausible that Plum is not morally responsible in Cases 1 and 2 are present in Case 3, then even if it is not as initially obvious that he is not responsible in this case, it seems we are nevertheless committed to thinking it in light of our judgments about the initial cases. Either that, or we owe incompatibilists like Pereboom an explanation of precisely how

the differences between Cases 1 and 2 and Case 3 *make a difference* for morally responsible agency.

But lacking that, we should also conclude that Plum is not morally responsible in Case 4. Of course, denying that Plum is responsible in this case *just is* a denial of compatibilism. It is true, of course, that Plum is not determined as a result of the interventions or programming or training of other agents in Case 4 (as he was in Cases 1–3). And yet this truth does not seem to bear on the fact that his decision to kill is ultimately explained by forces outside of his control. And since it is this feature of Cases 1–3 that Pereboom identifies as the most plausible explanation for why Plum is not responsible in those cases, we can reasonably adduce that contrary to compatibilists' protestations, he is not responsible in this case. In essence, the explanation for why Plum is not responsible in these cases is that he is not the source of his actions—that they originate not in his will or in *him* in any meaningful sense but in factors over which he bears no responsibility.

4. A Softer, Gentler Compatibilism

How might compatibilists reply to this challenge, raised first by Wisdom and precisified to a high degree in Mele's Zygote Argument and Pereboom's Four-Case Argument?⁸ One general strategy is to reject the assessment of Wisdom, Mele, and Pereboom that agents manipulated or programmed or "set up" in the ways considered are exempt from responsibility.⁹ Compatibilists who adopt this strategy are known as *hard* compatibilists. A second strategy has been to argue that there are morally significant differences between the respective machinations of Diana or the team of nefarious neuroscientists and causal determination by blind physical processes.¹⁰ In contrast to their hard compatibilist counterparts, philosophers who adopt this approach are known as *soft* compatibilists.

Each of these compatibilist strategies has something going for it, but in the remainder of this paper I want to offer some support on behalf of soft compatibilism. In particular, I will argue that *mere* causal determinism is in an important sense *irrelevant* to the behavior of determined agents. However, we will see, manipulation of the sort that Ernie and Plum undergo is clearly relevant. So the mere fact that an agent is causally determined does not bear on her status as morally responsible in the way that manipulation does. As a result, if compatibilism is false, it is not false because it entails that agents like Ernie or Plum are no different than "ordinary" agents who are causally determined to act as they do.

To arrive at these striking conclusions we have to temporarily step away from issues relating to compatibilism, manipulation, and determinism. Instead we will need to consider some rather mundane cases and investigate the ways

in which factors extrinsic to an agent's will affect her status as a morally responsible agent. We can then apply those lessons to the fantastical scenarios of the sort we see in Wisdom, Mele, and Pereboom.

5. The Importance of Being Independent

With this in mind, I want to consider the following two cases with an eye towards the question, "Does the degree to which these agents are praiseworthy differ?"

A1. James is a considerate partner, and sometimes—just because he knows it makes her happy—he buys his partner Pearl a bouquet of flowers. Today, as he approaches the florist on his way home from work, it pops into his head that he should stop in and buy some flowers for Pearl. Yet, just as he turns to go into the shop, she calls him. During the conversation she tells him that she had a rough day. She also mentions that if he did something for her, the otherwise terrible day would be, in some small way, salvaged. James does not mention to Pearl that he was already planning on buying her flowers, and when he brings the bouquet home, she is very grateful.

Before continuing, fix in your mind how praiseworthy (if at all) you take James to be in A1. Once you have that worked out, consider A2.

A2. James is a considerate partner, and sometimes—just because he knows it makes her happy—he buys his partner Pearl a bouquet of flowers. Today, as he approaches the florist on his way home from work, James is so tired from work that it does not even occur to him, as it sometimes does, to stop in at the florist and buy some flowers for Pearl. Yet, just as he passes the shop, she calls him. During the conversation she tells him that she had a rough day. She also mentions that if he did something for her, the otherwise terrible day would be, in some small way, salvaged. James forgets about his tough day, starts thinking about how he can make Pearl's day better, and decides to stop at the florist. When he brings the bouquet home, she is very grateful.

Again, before you go on, make a judgment of how praiseworthy (if at all) you take James to be in A2. Once you have made such a judgment, let us compare these cases more carefully.

To begin, let us note that in each of these cases, James brings Pearl a bouquet of flowers. So at that level of description he has performed the same, considerate action. Let us also note that in neither of these cases is he morally *obligated* to do so. Because James goes above and beyond his obligations in acting so considerately, it seems plausible that he genuinely *deserves* Pearl's gratitude in each case and so is praiseworthy (to some degree or other) for his actions.

Yet despite the fact that James seems to be praiseworthy (to some degree or other) in each scenario, I am inclined to think that James is *more praiseworthy* in A1 than he would be in A2. The reason for this, I will ultimately claim,

is that in A1 his decision to bring Pearl flowers is clearly *independent* from her request, since he would have bought Pearl flowers even if she had not called him and expressed that she was hoping to be cheered up. In light of this independence, it seems that James's subsequent action redounds more fully to him and so renders him more praiseworthy in A1 than in A2.

This explanation of the difference in the degrees to which James is praiseworthy in A1 and A2 needs to be spelled out a great deal before we will be able to see how it bears on cases of manipulation. But before I do so, I want to consider a parallel set of cases to see if the same intuition holds in cases of apparently blameworthy action, again with an eye towards the question, "Does the degree to which these agents are blameworthy differ?"

B1. James is a neighborhood tough. He regularly breaks into cars and steals any valuable electronics that he finds. Today, as he is walking down the street, he notices an expensive car. The owner is not around, and he thinks—as he sometimes does when he sees expensive cars—that he should break into it. Just after he makes this decision, his buddy Alex approaches from the other direction, also eyeing the car. Before James can say anything, Alex suggests that they break into the car and steal its electronics. James does not mention to Alex that he was already planning on doing so, and together, James and Alex break into the expensive car.

Again, make a judgment about the degree to which James is blameworthy (if at all) in B1 before continuing on to B2. Once you have that in your mind, consider the following:

B2. James is a neighborhood tough. He regularly breaks into cars and steals any valuable electronics that he finds. Today, as he is walking along the street, he notices an expensive car. The owner is not around, and he thinks—as he sometimes does when he sees expensive cars—that he should stop to take a look at it and admire it (James appreciates good cars). Just after he makes the decision to stop and admire it, his buddy Alex approaches from the other direction, also eyeing the car. Before James can say anything, Alex suggests that they break into the car and steal its electronics. Although James had not been considering doing so, he sees Alex's suggestion as attractive, and together, James and Alex break into the expensive car.

To what degree is James blameworthy (if at all) in this scenario? If you are like me, there will be some difference between your answer to this question and the one posed upon reflection on B1.

As we saw with A1 and A2, there are some important similarities between B1 and B2. In each of these cases, for example, it is true that James breaks into a car. And in each of these cases, it is reasonable to think that he is blameworthy for doing so. Indeed, it even seems reasonable to think that he is *sufficiently* blameworthy in each of these cases to warrant severe responses (e.g., criminal punishment). However, it also seems to me that James is *more* blameworthy in B1 than he is in B2. Again, a natural explanation for this intuition is that in B1

James's decision does not depend on Alex's malevolent suggestion. If Alex had not been waltzing down the same street (at the exact moment that James passed the car), James would not have broken into the car in B2. But this is not the case in B1. Even if Alex had not been present in B1, James would have broken into the car.

So again, we have two cases in which it seems differences in the degree to which an agent's action depends on factors extrinsic to her will are relevant to the degree to which he or she is morally responsible for the action. But what exactly does this mean? And why should the differences described here matter?

As I have said, it seems that in A1 and A2 James is praiseworthy for his considerate deed. And similarly, it seems that in B1 and B2 James is blameworthy for his malicious deed. But in A1 and B1, James seems to be *more* praiseworthy and blameworthy (respectively) for his actions than he is in A2 and B2. The reason for this, I submit, is that in A2 and B2, had Pearl and Alex not intervened, James would not have bought the flowers or broken into the car. If this is a plausible explanation for why degrees of responsibility diverge in these cases, then it suggests something more general: the degree to which an agent is praiseworthy or blameworthy for her action depends (in part) on what they would have done had their action been *independent* from the other agents' interventions.¹¹

Independence of this sort matters because it seems to reliably track the fact that the act in question was the agent's own. A natural corollary to this is the idea that the degree to which my act is independent of factors extrinsic to my will, determines the degree to which the act is really mine in some robust sense—i.e., the degree to which it truly redounds to *me* and not to something or someone else.

6. Manipulation and Independence

If this is correct, then there is good reason to think that divinely designed agents like Ernie and neurologically manipulated agents like Plum are less blameworthy for their actions than those agents who are merely causally determined by blind physical processes. To develop this point in some detail, I will limit my focus in the remainder of the paper to Pereboom's Four-Case argument.¹²

To arrive at conclusions about the degree to which agents' actions are independent from factors extrinsic to their wills, we compared not the actions agents performed but whether they would have performed those actions themselves. So to assess the degree to which Plum's action is meaningfully *his*, we can ask whether there is any reason to think that he would have murdered White absent the interventions of the neuroscientists in Cases 1 and 2. I doubt there is. But if not, then the following seems true:

Case 4: If Plum were not neurologically manipulated to murder White, it would not have been the case that Plum would have murdered White.¹³

Of course, merely stating my doubts about what Plum would do in the absence of the neuroscientists is not itself an *argument*. So understanding (4) and why one might come to regard it as true will prove important for the compatibilist.

Why, then, accept (4)?

We should accept it, I think, because of what we know about the kinds of manipulators described in Cases 1 and 2 and standard counterfactual semantics.¹⁴ Consider that the nearest world in which the antecedent of (4) obtains plausibly is not a world in which the consequent obtains, either because Plum is not created, or because the neuroscientists have different aims and so manipulate him to act in some other way. Plum's behavior (and even his existence in Case 1) seem to hang precariously on the content of neuroscientists' intentions, whether those intentions lead to a world in which Plum does not exist or a world in which he is manipulated to act differently.¹⁵ In either case, the truth of (4) seems assured. Accordingly, we have grounds for concluding that Plum's decision to murder depends on neurological manipulation in some very strong sense.¹⁶

A strong sense of dependence also figures importantly in A2 and B2. Consider:

Case 5: If Pearl had not called James in A2, it would not have been the case that he bought her flowers.

In other words, in A2 James would not have bought the flowers if Pearl had not asked for them. The same can be said of James in B2.

Case 6: If James had not seen Alex in B2, it would not have been the case that James broke into the expensive car.

That is, in B2 James would not have broken into the car if Alex had not first suggested it.

Now let us recall that facts about the causal history of James's actions reflect a difference in the *degree* to which James was praise- and blameworthy. Similarly, we should expect that because the neuroscientists' role in the causal history of Plum's action supports the truth of a relevantly similar counterfactual in Cases 1 and 2, the degree to which Plum is responsible and blameworthy for his action is also mitigated. In other words, just as the strong degree of dependence present in (5) and (6) reflects something in the causal history of James's action that is relevant to the degree to which he is responsible, the presence of the same sort of dependence in (4) suggests that in two cases Pereboom gives of neurological manipulation, the degree to which an agent is responsible is affected.

But now we must ask: would causal determinism figure in the etiology of

agents' actions in the same way? And if not, then there is a normatively significant difference between manipulation of the sort present in Cases 1 and 2 and causal determinism as it figures in Case 4.

7. The Irrelevance of Determinism

Concerning cases of causal determination by blind physical processes, I think there are good grounds for thinking that Plum's decision to murder White is importantly *independent* from the fact that he was causally determined (by blind physical processes) to make that decision. Chief among these is that in Case 4 the fact that Plum is causally determined to murder White cannot bear on the actual-sequence etiology of his decision to do so. If it did, then it would be true of Plum that had he not been causally determined to murder White, he would not have done so. But this is not the case. After all, in the closest world or set of worlds in which causal *indeterminism* obtains (i.e., the closest world in which he is not causally determined to murder White), it is plausible to think that the same prior events would still have occurred and indeterministically issued in Plum's murdering White, since he would still reason egoistically, it would still be the case that murdering White would further his aims, etc. Indeed, in this set of worlds, the relata in the causal chain that leads to White's murder would be no different then if determinism were true.¹⁷ In other words, the mere introduction of causal indeterminism to the actual-sequence etiology of Plum's decision is not, by itself, relevant to whether a type-identical causal sequence in fact obtains.¹⁸

Of course, the evaluation of such counterfactuals is admittedly delicate and contentious. It is nevertheless plausible to think that:

Case 7: If Plum were not causally determined (by blind physical processes) to murder White, it would have nevertheless been the case that Plum would have murdered White.

As I see it, (7) simply codifies the intuition that even if the causal sequence leading up to his decision had been indeterministic, Plum would have chosen, for the same egoistic reasons that drove his choice in Case 4, to murder White.¹⁹ But notice, something like (7) is also true of James in the situation described in cases A1 and B1. Consider, for example:

Case 8: If Pearl had not called James in A1, it would have been the case that he bought her flowers.

And likewise, consider (9):

Case 9: If James had not seen Alex in B1, it would have been the case that James broke into the car.

As we saw, James was more praiseworthy and blameworthy (respectively) for his actions in A1 and B1 than he was in the alternative scenarios of A2 and B2. If this is true, then we should think that Plum is more blameworthy for murdering White assuming causal determination by blind physical processes (i.e., Case 4) than he is as described by Pereboom in Cases 1 and 2 (i.e., cases of neurological manipulation). After all, the difference between (4) and (7) is the same as the difference between (5) and (8), and as we saw from the cases involving James, that difference, while apparently only minor, is relevant to James's moral responsibility because it is relevant to the degree to which it is appropriate for Pearl to be grateful for the flowers. Likewise, the difference between (4) and (7) is the same as the difference between (6) and (9). And as we saw, that difference is similarly relevant to James's moral responsibility because it is relevant to the degree to which it is appropriate for us to resent or otherwise blame James for breaking into the car. So given (7) it seems reasonable to conclude that in Case 4, Plum's action redounds to him in a significant way. And when actions redound to us in this way—when they reflect what we would do given the absence of all sorts of irrelevant features (like causal determinism)—then it is quite plausible that these are precisely the class of actions for which we are responsible.

8. Conclusion

Thus, there is an important difference between the neurological manipulation as described in Cases 1 and 2 and mere causal determination by blind physical processes as described in Case 4. And further, as I argued earlier, this difference is present in at least some ordinary cases in which we alter judgments of praise and blame, gratitude and resentment. So, the difference between Cases 1 and 2 on the one hand and Case 4 on the other is, as they say, a difference that *makes* a difference. While the actions of causally determined agents are in some real sense independent of the causal determinism itself, the actions of the manipulated agents are not independent of the manipulation—at least if the manipulation looks anything like the sort described by Wisdom, Mele, or Pereboom. Compatibilism might not be true, but if it is false, it is not because it is too much like manipulation.

Author Note

I'd like to acknowledge Neal Tognazzini, Philip Swenson, Patrick Todd, and John Fischer for helping me think through these ideas.

1. He also reports anecdotally at least that he was perhaps the first “experimental philosopher” in the contemporary sense of that term when he notes that he has “confirmed what [he] seem[s] to see from inspection by asking others to inspect the same problem” (Wisdom 1934, 116).

2. The most prominent defense of this claim is found in Peter van Inwagen (1983).

3. An important instance of this line of thought is David Lewis (1981).

4. I say “purported” because many have denied that it does this. For a (partial) sampling, see van Inwagen (1978), David Widerker (1995), Laura Ekstrom (1999), Daniel Speak (2007), and Christopher Evan Franklin (2011). For defenses of Frankfurt, see Fischer (1982, 2010), Derk Pereboom (2001), and Filipe Leon and Neal A. Tognazzini (2010).

5. Fischer does not actually cite Wisdom (1934) and given the paucity of citations to that text in the other articles at the time, it is doubtful that its insights—significant as they are—played any causal role in history of the views developed by Fischer here or later in Al Mele (1995, 2006) or Derk Pereboom (1995, 2001).

6. Mele himself is agnostic about the soundness of this argument. He does, however, regard it as a serious challenge to compatibilism that cannot be ignored.

7. Plum does not satisfy every possible compatibilist condition as written, but this is not a failing of case, since Pereboom captures a number of leading compatibilist views. So for any possible compatibilist condition that is not mentioned, it is the compatibilist’s burden to show that (1) that condition is met by agents in deterministic universes but not by Plum, and (2) is well-motivated (i.e., is not an ad hoc anti-manipulation condition).

8. There are, of course, further ways that incompatibilists have sought to shore up something like Wisdom’s thought. Excepting Mele and Pereboom, perhaps the most sophisticated versions of these arguments are due to Patrick Todd (2011, 2013).

9. See Gary Watson (1999), Michael McKenna (2007), John Martin Fischer (2011), and Daniel Haas (2013).

10. See Kristin Mickelson (2010), Robyn Repko Waller (2013), and Oisín Deery and Eddy Nahmias (2017). See also Martin Montminy and Daniel Tinney (2018) for a reply that has elements of hard and soft compatibilism.

11. For more on the significance of interventions see Deery and Mahmias (2017).

12. What I say in the rest of the paper is meant to apply *mutatis mutandis* to Mele’s Zygote Argument.

13. Here I frame things in terms of Pereboom’s Four-Case Argument, but what I say in the remainder of the paper will apply *mutatis mutandis* to Mele’s Zygote argument.

14. See David Lewis (1973).

15. For an earlier attempt to understand the significance of manipulators’ intentions, see Waller (2013).

16. Note that I not reductively analyzing the notoriously vexed notion of dependence in terms of counterfactual dependence here (such reduction is undoubtedly doomed to failure). However, I am suggesting that (at least in this kind of case) the counterfactual dependence of Plum’s decision to murder White on the fact that he is neurologically manipulated gives us good reason to think that this decision depends on the neurological manipulation in some deeper sense. Thus on this picture, counterfactual dependence of the sort expressed by (1) is a manifestation of a deeper, perhaps unanalyzable kind of dependence.

17. Though in such a case, this causal chain would be such that the probability of the causal antecedents producing the relevant outcome (i.e., Plum murdering White)

would be less than 1.0 (as would be the case given the truth of causal determinism). Of course, the nearest world to the world in which Plum is causally determined, the relevant probabilities would undoubtedly be high—perhaps .999. And it is in light of this fact that we have good reason to accept (2), which I argue for below.

18. Of course, the introduction of causal indeterminism does affect whether the antecedent causes are sufficient for Plum's murdering White, and this might affect whether Plum has the ability to do otherwise (it certainly affects whether holding fixed the past and the laws, Plum has the ability to do otherwise). But this is not germane to the issue at hand.

19. Of course, the mere fact that on the assumption of causal indeterminism it is true that Plum would choose to murder White does not undermine the further fact that some incompatibilists (though not Pereboom) take to be paramount for responsibility—namely, that Plum could have chosen to do otherwise, even holding fixed the past and the laws of nature.

Works Cited

- Deery, Oisín and Eddy Nahmias. "Defeating Manipulation Arguments: Interventionist Causation and Compatibilist Sourcehood." *Philosophical Studies* 174.5. 2017.
- Ekstrom, Laura. *Free Will: A Philosophical Study*. Westview. 1999.
- Fischer, John Martin. "Responsibility and Control." *Journal of Philosophy* 79. 1982.
- _____. "The Frankfurt Cases: The Moral of the Stories." *Philosophical Review* 119.3. 2010.
- _____. "The Zygote Argument Remixed." *Analysis*, 2011.
- Franklin, Christopher. "Neo-Frankfurtians and Buffer Cases: The New Challenge to the Principle of Alternative Possibilities." *Philosophical Studies* 152. 2011.
- Haas, Daniel. "In Defense of Hard-line Replies to the Multiple-Case Manipulation Argument." *Philosophical Studies* 163.3. 2013.
- Leon, Felipe, and Neal A. Tognazzini. "Why Frankfurt-Examples don't need to Succeed to Succeed." *Philosophy and Phenomenological Research*. 80.3. 2010.
- Lewis, David. *Counterfactuals*. Malden: Wiley-Blackwell. 1973.
- McKenna, Michael. "A Hard-line Reply to Pereboom's Four-case Argument." *Philosophy and Phenomenological Research* 77. 2008.
- Mele, Al. *Autonomous Agents*. New York: Oxford University Press. 1995.
- _____. *Free Will and Luck*. New York: Oxford University Press. 2006.
- Mickelson, Kristin. "The Soft-Line Solution to Pereboom's Four-Case Argument." *Australasian Journal of Philosophy* 88.4. 2010.
- Montminy, Martin, and Daniel Tinney. "Manipulation and Degrees of Blameworthiness." *Journal of Ethics*. 2018.
- Pereboom, Derk. "Determinism Al Dente." *Nous* 29.1. 1995.
- _____. *Living Without Free Will*. Cambridge: Cambridge University Press. 2001.
- _____. *Free Will, Agency, and Meaning in Life*. Oxford: Oxford University Press. 2014.
- Speak, Daniel. "The Impertinence of Frankfurt-Style Argument." *Philosophical Quarterly* 57.226. 2007.
- Todd, Patrick. "A New Approach to Manipulation Arguments." *Philosophical Studies* 152.1. 2011.
- _____. "Defending (a Modified Version of) the Zygote Argument." *Philosophical Studies* 164.1. 2013.
- van Inwagen, Peter. "Ability and Responsibility." *Philosophical Review* 87.2. 1978.

- _____. *An Essay on Free Will*. Oxford: Oxford University Press, 1983.
- Waller, Robyn Repko. "The Threat of Effective Intentions to Moral Responsibility in the Zygote Argument." *Philosophia* 43.4. 2013.
- Watson, Gary. "Soft Libertarianism and Hard Compatibilism." *Journal of Ethics* 3. 1999.
- Widerker, David. "Libertarianism and Frankfurt's Attack on the Principle of Alternative Possibilities." *Philosophical Review* 104.2. 1995.
- Wisdom, John. *Problems of Mind and Matter*. Cambridge: Cambridge University Press. 1934.
-

D. Justin Coates is an associate professor of philosophy at the University of Houston. He primarily works on issues related to ethics and the philosophy of action. He is the co-editor of *Blame: Its Nature and Norms* (2013), and *Oxford Studies in Agency and Responsibility, Vol. 5: Essays on themes from the work of Gary Watson* (2019), and his work has appeared in journals such as *Philosophical Studies*, *American Philosophical Quarterly*, *Canadian Journal of Philosophy*, and *British Journal for the History of Philosophy*. Department of Philosophy, Agnes Arnold Hall 513, University of Houston, 4800 Calhoun Rd., Houston, TX 77004. <djcoates@uh.edu>.