

Why History Matters for Moral Responsibility: Evaluating History-Sensitive Structuralism

By Taylor W. Cyr

Forthcoming in *Philosophical Issues*; please cite published version.

Abstract Is moral responsibility essentially historical, or does an agent's moral responsibility for an action depend only on their psychological structure at that time? In previous work, I have argued that the two main (non-skeptical) views on moral responsibility and agents' histories—historicism and standard structuralism—are vulnerable to objections that are avoided by a third option, namely history-sensitive structuralism. In this paper, I develop this view in greater detail and evaluate the view by comparing it with its three dialectical rivals: skepticism about moral responsibility, historicism, and standard structuralism. Each comparison includes discussion of new work on moral responsibility and agents' histories, and along the way I offer new arguments for preferring history-sensitive structuralism, paying special attention to the view's explanatory power.

Keywords historicism, luck, manipulation, moral responsibility, skepticism, structuralism

1. Introduction: Structuralism, Historicism, and a Middle Path

Structuralism is the view that whether an agent is morally responsible for an action depends only on their psychological structure at that time.¹ One of the main challenges for structuralists comes from cases of manipulation. Consider the following case (adapted from Mele 2006: 164-166):

Ann is a free agent and an exceptionally industrious philosopher. She puts in twelve solid hours a day, seven days a week, and she enjoys almost every minute of it. Beth, an equally talented colleague, values many things above philosophy for reasons that she has refined and endorsed on the basis of careful critical reflection over many years. Beth identifies with and enjoys her own way of life, and she is confident that it has a breadth, depth, and richness that long days in the office would destroy. Their dean wants Beth to

¹ I am restricting my focus to *direct* (or *non-derivative*) moral responsibility. Without the restriction, no one who thinks indirect moral responsibility can be indirect (as in cases of “tracing”) would count as a structuralist (cf. McKenna 2012: 156). Structuralists include Brink (2021), Frankfurt (1988), Vargas (2013), and Watson (2004). In this paper, I am concerned only with direct (or non-derivative) moral responsibility.

be like Ann. Without the knowledge of either philosopher, he hires a team of psychologists to determine what makes Ann tick and a team of new-wave brainwashers to make Beth like Ann. The psychologists decide that Ann’s peculiar hierarchy of values accounts for her productivity, and the brainwashers instill the same hierarchy in Beth while eradicating all competing values—via new-wave brainwashing, of course. Beth is now, in the relevant respect, a “psychological twin” of Ann. After Beth is brainwashed, she does some extra philosophical work that she would not have done had she not been manipulated—she stays at the office into the evening to review a manuscript for a journal.

Given that she was manipulated into performing the action, Beth seems not morally responsible for reviewing the manuscript; however, given that she is a psychological twin of Ann, Beth satisfies any merely structural conditions on moral responsibility for action that the structuralist might propose. These considerations lead some philosophers to reject structuralism and to endorse historicism instead, maintaining that an agent’s history (how the agent came to have the particular psychological structure from which they act) can make a difference as to whether or not they are morally responsible for an action.²

In previous work (2020), I have argued against historicism as well as against standard versions of structuralism, and my objections to both sets of views stem from considerations pertaining to constitutive luck. Constitutive luck is luck in being the kind of person you are, or in having the character that you have (cf. Nagel 1979). Unless one endorses skepticism about moral responsibility, one should admit the possibility of agents being morally responsible despite being entirely constitutively lucky—at the very least, at the beginning of agents’ careers as morally

² Historicists include Mele (1995; 2006; 2019) and Fischer and Ravizza (1998).

responsible agents. But this leads to a problem for historicism, which can be formulated as follows:

1. Agents who are entirely constitutively lucky can be morally responsible for what they do.
2. There is no relevant difference between agents who have been manipulated in certain ways (as in Beth's case) and agents who are morally responsible for actions that stem from characters with respect to which they are entirely constitutively lucky.
3. Thus, manipulated agents (like Beth) can be morally responsible for what they do, despite failing to satisfy historical conditions on moral responsibility, which is to say that historicism is false.³

The problem for standard structuralism, by contrast, is that it implies that there is no moral responsibility-relevant difference between Ann and Beth with respect to their moral responsibility for the action in question. The standard structuralist will say that Ann, but not Beth, is morally responsible for having the character from which she acts, which is some moral responsibility-relevant difference between the two, but even so the standard structuralist does not take seriously the way in which constitutive luck *mitigates* moral responsibility for actions; for example, if Beth is morally responsible for reviewing the manuscript at all, she seems *less* morally responsible for doing so than does Ann.⁴

³ I do not mean to suggest that the falsity of historicism (according to which an agent's history *can* make a difference to whether they are morally responsible) follows just from the fact that *Beth* is morally responsible. Rather, my aim is to argue for the falsity of historical conditions on moral responsibility directly, and what I say about the case of Beth could be modified to any sort of case proposed in support of a historical condition. Thanks to Carolina Sartorio for raising this point.

⁴ Other cases provide further support, including McKenna's (2004) case of Suzie Instant and my (2020) case of two little agents, Little Tony and Little Anthony.

After raising these challenges, I sketched an alternative, *history-sensitive structuralism*, according to which an agent's history cannot (by itself) undermine their moral responsibility for what they do, but it can affect the degree to which they are morally responsible for what they do. In this paper, I develop this view in greater detail (section 2) and evaluate the view by comparing it with its three dialectical rivals: skepticism about moral responsibility (section 3), historicism (section 4), and standard structuralism (section 5). Each comparison includes discussion of new work on moral responsibility and agents' histories, and along the way I offer new arguments for preferring history-sensitive structuralism, paying special attention to the view's explanatory power.

2. *History-Sensitive Structuralism: Filling in the Details*

History-sensitive structuralism is a version of structuralism since it holds that structural features of an agent are sufficient to establish that they meet the threshold conditions on moral responsibility (contra historicism). Unlike standard structuralist positions, however, history-sensitive structuralism allows for an agent's history to affect their moral responsibility for what they do. Consider again the case of Ann and Beth. Whereas historicists typically take Beth's moral responsibility for reviewing the manuscript to be undermined by the manipulation, the history-sensitive structuralist maintains that Beth is morally responsible for that action. Given that Beth has been manipulated, though, making (the relevant part of) her constitution at the time of that action a matter of luck for her, the history-sensitive structuralist will say that Beth is *less* morally responsible for that action than Ann is for her action.

Crucial to this sketch are two presuppositions: first, that moral responsibility comes in degrees; and, second, that an agent's degree of moral responsibility for an action may depend on their degree of constitutive luck. I find the first of these claims is intuitively plausible, and I am

not alone in holding this view, though others take moral responsibility to be a binary notion and something related—typically praiseworthiness and blameworthiness—to be scalar.⁵ I will assume that moral responsibility itself is scalar, but the view developed here could be adapted, *mutatis mutandis*, to fit with the alternative view. The second claim noted above (that an agent’s degree of moral responsibility for an action may depend on their degree of constitutive luck) merits further consideration, which I will give it in the remainder of this section. I will start by developing an account of how an agent’s constitutive luck may affect their degree of moral responsibility for action, and I will then develop an account of how an agent’s constitutive luck may be mitigated over time—which will explain how an agent’s degree of moral responsibility for actions may increase over time.

My account of how an agent’s constitutive luck may affect their degree of moral responsibility for action is based on an idea I floated in previous work (2019). The idea is that the degree to which agents are morally responsible is a function of the degree to which they meet the various conditions on moral responsibility, whatever those conditions may be. Suppose that we follow tradition and adapt Aristotle’s two conditions on voluntary action (*Nicomachean Ethics* 1109b30–1111b5), one concerning lack of ignorance and the other concerning lack of compulsion, into the *epistemic* and *control* (or *freedom*) conditions on moral responsibility.⁶ While there would, of course, be some threshold of knowledge (or perhaps merely justified belief) and control required in order for an agent to be morally responsible for some behavior, my view is that the exact amount of knowledge/control had *above* that threshold determines the

⁵ For some discussion, see Coates and Swenson (2013) and Nelkin (2016).

⁶ For more on the *Aristotelian conditions*, as they are often called, see Fischer and Ravizza (1998: 12–14).

agent's degree of moral responsibility for that behavior. In what follows, I focus on the control condition.⁷

There is much to be said for the view that a person's degree of control plays a role in determining their degree of moral responsibility. First, it is clear that control comes in degrees, not simply for our overt actions and their consequences (e.g., an expert archer has more control over whether their arrow hits the target than I do) but also for our *basic* mental actions of deciding (e.g., I have more control over whether I decide to refrain from eating the cookie when I am well-rested and not under stress). Moreover, we also factor in a person's degree of control over their behavior when considering whether to respond with praise or blame *and how much* praise or blame with which to respond, as when we take duress (a diminishment of control) to affect what sort of responsibility practice is appropriate, and to what extent.⁸ Further support comes from what I think of as the *villain backstory* phenomenon, familiar from literature and film, which happens when we learn about the history of a character whom we are initially poised to blame and come to temper our reactions in light of the character's mitigated control.⁹

Supposing that an agent's degree of moral responsibility tracks their degree of control (perhaps among factors), it is easy to see why an agent's constitutive luck may affect their degree of moral responsibility for action. After all, constitutive luck is a form of luck, and luck and

⁷ The reason for this is that luck (including constitutive luck) is inimical to *control*—perhaps the two are even inversely related—and so it is this condition that will be relevant here even if there is a distinct epistemic condition on moral responsibility. For some reasons to be skeptical that there is a distinct epistemic condition on moral responsibility, see Mele (2010; Forthcoming).

⁸ An alternative to this picture, which I do not have space to consider in detail here, is the view that the agent's degree of moral responsibility is not a function of the agent's degree of control but rather of the agent's quality of will. After all, if two agents manifest the same degree of ill will in performing the same wrong act, they may seem morally responsible to the same degree even if one agent exercised a greater degree of control in performing the wrong act. Thanks to Michael McKenna for raising this point.

⁹ See also Watson's (1987) discussion of Robert Harris.

control are inversely related: the more that something is under one's control, the less it is a matter of luck for them; and the more that something is a matter of luck for a person, the less it is under their control. When a person's actions stem from a character or set of values with respect to which they are lucky (lacked control over), that person's actions are less under their control than they would be had their character and values been something under their control. To the extent that an agent is constitutively lucky, then, and all else being equal, the agent satisfies the control condition on moral responsibility to a lesser degree—and thus, on my view, is less morally responsible than they would be if they were less constitutively lucky.¹⁰

That, on my view, is how an agent's constitutive luck may affect their degree of moral responsibility for action, and we can now turn to the question of how an agent's constitutive luck may be mitigated over time. While we do not (at least typically) have direct, voluntary control over our character, values, etc., how we are constituted is often shaped by the choices that we make (among other things), and in such cases we indirectly control our constitution at subsequent times. Returning to the case of Ann and Beth, we have standard example of how this works. Before she is manipulated, Beth values many things above philosophy for reasons that she has refined and endorsed on the basis of careful critical reflection over many years, and so those values, though perhaps a matter of luck for her before the process of refinement and endorsement, are not completely a matter of luck for her anymore. And this, in general, is the process by which young children, who are never morally responsible for what they do, grow into

¹⁰ The *ceteris paribus* clause is needed here because it could be the case that an agent's being more constitutively lucky coincides with their being less lucky overall—perhaps they are a lot less circumstantially lucky. Though my focus here is constitutive luck, I will mention other forms of luck below, and I will here set aside the question of how to determine an agent's degree of control when different forms of luck pull in opposite directions. Thanks to Carolina Sartorio for this point.

older children and adults, who are sometimes morally responsible for what they do—with degree of moral responsibility increasing along the way.¹¹

My account charts a middle path between the view that we create ourselves *ex nihilo*, on the one hand, and the view that constitutive luck is ineliminable, on the other. My account sees the mitigation of constitutive luck as the process of self-*re*-creation, to borrow a term from Joel Feinberg, who says:

A common-sense account of self-creation...can be given, provided we avoid the mistake of thinking that there can be no self-determination unless the self that does the determining is already formed...The extent of the child's role in his own shaping is...a process of continuous growth already begun at birth... (Feinberg 1986: 33-34, quoted in Fischer 2012: 165)

When Beth is manipulated into having the psychological profile of Ann, she is re-created, to an extent, but not by herself. When the action of reviewing the manuscript results from Beth's new constitution, Beth has less control over that action than does Ann over her parallel action. Beth is more like a young child who is endowed with a character and set of values over which they had no control but which, with the right opportunities, could be reflected on, shaped, and endorsed, but only with the time required for such a process.

3. History-Sensitive Structuralism vs. Skepticism

The account I have just developed is partly shaped by worries raised by such skeptics about moral responsibility as Galen Strawson (1994) and Neil Levy (2011), both of whom argue that the pervasiveness of luck prevents the emergence of moral responsibility in the first place. History-sensitive structuralism is partly motivated by these concerns, insofar as it acknowledges

¹¹ See Cyr (2021) and Hartman (2023) for more examples, as well as for an argument that this process itself requires other forms of luck, namely resultant and circumstantial.

luck's mitigating effects on moral responsibility, but the view also "stares down" luck, insofar as it accepts that an agent can be morally responsible despite action from a constitution with respect to which they are entirely constitutively lucky, thereby avoiding a skeptical conclusion about moral responsibility.¹²

More recently, though she does not explicitly endorse a total skepticism about moral responsibility, Marcela Herdova (2020) has introduced a new skeptical worry. Drawing from L. A. Paul's (2014) influential work on transformative experiences, Herdova calls into question how much our characters and values are ever up to us, posing a challenge for views like mine according to which our constitutive luck may be mitigated over time. Herdova takes the challenge to strengthen the case for moral responsibility skepticism. After recounting her argument, I will respond by arguing that history-sensitive structuralism can explain the effects of transformative experiences in a way that rival accounts (historicism and standard structuralism) cannot and that the case for skepticism is not ultimately strengthened.

According to Herdova, the following claim captures a common thought but turns out to be false:

New Leaf Hypothesis (NLH): We often rationally and intentionally shape our characters and values in major ways (2020: 279)

Herdova presents two arguments against NLH, noting that the second may be a reformulation of the first (2020: 283-288). The two key claims in Herdova's case against NLH are the following: first, if an agent's traits or values undergo a major change, then this change constitutes a transformative experience (cf. Paul 2014); and, second, given the nature of transformative experiences, an agent undergoing such a change *cannot rationally and intentionally* (choose to)

¹² See Cyr (2019) and Hartman (2018) for responses to Levy and Strawson, respectively.

gain their new values or traits for the reasons that might justify such an action.¹³ While I would push back against some of the moves Herdova makes in support of these key claims, I will simply concede for the sake of argument that NLH is false, since my concern here is what Herdova goes on to say about why the falsity of NLH would be relevant to moral responsibility.

Herdova says that the falsity of NLH is relevant to moral responsibility for several reasons, two of which I will discuss here. The first is that, on her view, the falsity of NLH makes the case for skepticism about moral responsibility stronger:

If NLH is false, this severely limits the ways in which an optimist [a non-skeptic about moral responsibility] may rebut the claim that our characters and values are largely not up to us. What's more, the kinds of reasons that support the falsity of NLH provide, in and of themselves, *additional* evidence for thinking that our characters and values are not up to us in any significant way. (2020: 291)

The second reason Herdova thinks that the falsity of NLH is relevant to moral responsibility is because it turns out, on her view, that transformative experiences are not relevantly different from manipulation cases: “If manipulation cases show that certain agents are not responsible (or are less responsible) for their actions, then, given the falsity of NLH, any *major* changes in character and values that an agent undergoes also undermine her responsibility in a similar way” (2020: 291).

In response, and taking Herdova's second point first, I am inclined to agree. That is, I think Herdova is right to point out that transformative experiences are like manipulation cases

¹³ The first of these claims may be an incomplete characterization of transformative experiences (thanks to Carolina Sartorio for pointing this out to me), but for my purposes here it will be fine simply to follow Herdova's characterization of transformative experiences and to argue that such transformative experiences are relevantly like certain manipulation cases (and should be treated in a parallel way).

with respect to moral responsibility. But I think that this is exactly what we should expect, if history-sensitive structuralism is correct, and the parallel between transformative experiences and manipulation cases is better explained by my view than by historicism or standard structuralism. Consider the transformative experience of a religious convert (cf. Arpaly 2003: 127). Let us suppose that, before the conversion, our convert was a despicable person who was morally responsible for his despicable character.¹⁴ Upon “seeing the light” in a way that counts as a genuine transformative experience, our convert finds in himself new values and desires and performs a morally good action because of his improved character. On my view, the convert may be morally responsible for the morally good action he performs, but he will be significantly *less* morally responsible for doing so than would be some person who had exercised control over the formation of such a character, just as a manipulated agent like Beth would have mitigated moral responsibility. This strikes me as exactly the right thing to say, yet historicists and standard structuralists cannot say this. As Herdova argues, the historicist will have to say that, if there’s no relevant difference between transformative experiences and manipulation cases, then the convert is not even a little bit morally responsible for his good deed, and so the scope of moral responsibility is decreased, which may be very significant depending on how frequent transformative experiences turn out to be. But the standard structuralist will have to say that the convert is just as morally responsible for the good deed as would be someone who performed the same action from a character that they had shaped.

What I have just said in response to Herdova’s second point about the relevance of the falsity of NLH for moral responsibility is also relevant to her first point. Herdova thinks that, if NLH is false, this makes the case for skepticism about moral responsibility stronger, and perhaps

¹⁴ See, for example, Mele’s case of Chuck in Mele (1995: 162-163; 2006: 171).

it would be if historicism were the only viable option. But the history-sensitive structuralist can resist any potential pull toward skepticism by maintaining both 1) that transformative experiences do not undermine moral responsibility, and 2) that the more minor changes (less than *major* changes) involved when an agent's constitutive luck is mitigated do not require the truth of NLH. Both of these are plausible, on history-sensitive structuralism, and taken together we have a reason for doubting Herdova's claim that falsity of NLH would strengthen the case for moral responsibility skepticism.

4. History-Sensitive Structuralism vs. Historicism

History-sensitive structuralism is a structuralist position, and so is not a historicist position, but it is partly motivated by the objections historicists have tended to raise against standard structuralism. As indicated above, historicists often appeal to cases of manipulation to call into question the structuralist's position that moral responsibility is non-historical. The problem I have raised for the historicist, however, is that reflection on those cases of manipulation, together with cases where agents are first beginning to be morally responsible, should lead us to think that the manipulated agents historicists bring up *are* in fact (a little bit) morally responsible for what they are manipulated to do. Recall the argument introduced above:

1. Agents who are entirely constitutively lucky can be morally responsible for what they do.
2. There is no relevant difference between agents who have been manipulated in certain ways (as in Beth's case) and agents who are morally responsible for actions that stem from characters with respect to which they are entirely constitutively lucky.
3. Thus, manipulated agents (like Beth) can be morally responsible for what they do, despite failing to satisfy historical conditions on moral responsibility, which is to say that historicism is false.

If we are to avoid skepticism, we must accept the first premise. So, in order to resist this conclusion, historicists will need to reject the second premise and point to some relevant difference between manipulated agents and agents who are first beginning to be morally responsible.

In recent work, Mele (2020; 2021) has done exactly this. He has defended historicism against my argument by rejecting the second premise and proposing a relevant difference between little agents and manipulated agents, namely that little agents typically have some control over the values from which they perform the first actions for which they are morally responsible, whereas manipulated agents do not have such control. Mele thinks that this difference is relevant to moral responsibility because he defends a principle, *NFMR* (2020: 3151; cf. Mele 2019: 66-67), which is a *sufficient* condition for an agent's *not* being morally responsible, and *NFMR* includes being manipulated in the sort of way that Beth was (2020: 3151ff.). An implication of that condition is that a *necessary* condition of an agent's being morally responsible is the agent's not having been manipulated, and little agents satisfy this necessary condition on moral responsibility while manipulated agents do not.

Supposing (as we are) that the little agents in question have not yet performed any action for which they were morally responsible, I do not think that the difference Mele identifies bears on the moral responsibility of the agents in question. In my earlier work, I noted that "if an agent is not morally responsible for bringing about a change to her character, then it does not make a difference to her constitutive luck whether her character was influenced by her own past actions (for which she was not morally responsible) or someone else's or no one's at all" (2020: 2387, n. 8). But Mele disagrees, elaborating as follows (where little Tony is an ordinary "little agent"):

Even before [little Tony] performed the first action for which he was morally responsible, he had some control over his value system (he was capable of reflecting on and learning from the consequences of his behavior and modifying his behavior accordingly, for example)—it just was not moral-responsibility-level control; the capacities at issue were not quite robust enough yet. (2021: 309)

If we were to compare little Tony with “little Anton” (cf. Cyr 2020: 2389), who is exactly like Tony except that he has just been manipulated into having the same hierarchy of values as Tony, Mele would have to say that Tony is morally responsible but that Anton is not.¹⁵ Although neither exercised “moral-responsibility-level control” over their values, Tony exercised some control over his in a way that Anton did not.

I do not think that this is the correct result. By stipulation, little Tony and little Anton have the same capacities, hierarchy of values, and every other structural feature we might point to, and neither is morally responsible for their having those structural features at the time of the action in question. Given this, it strikes me as intuitive that Anton is morally responsible too. Still, I suspect that historicists will report different intuitions, and it would be nice to do more than to register a different judgment about this case. So, in the remainder of this section, I want to sketch two arguments for preferring history-sensitive structuralism over historicism, both of which highlight explanatory advantages of the former.

The first is the argument from restored moral responsibility. Both Mele and I agree that manipulated agents can eventually become morally responsible for acting from their implanted values. But when, and how, does that happen? Mele does not give a full account of this, but he does tell an extended story about Beth in which, a year after her manipulation, she has come to

¹⁵ And Mele does say this about “little Tara” (2021: 309), a case from McKenna (2021) that is not relevantly different from the case of little Anton.

“reflectively embrace and identify with the values at issue—products of manipulation—as part of a package that supports what she always valued most” (2020: 3149), where what she has always valued most was unaffected by the manipulation. It is unclear why this cannot happen immediately upon Beth’s manipulation, though Mele does say that she “benefits from a year’s worth of experience of a life organized around hard philosophical work (2020: 3149), and perhaps that is crucial to the embrace of the implanted value.

But the history-sensitive structuralist has the resources to explain both when and how moral responsibility is restored. It happens right away (so perhaps it does not need to be *restored*, strictly speaking), since the fact that the agent was manipulated into acting from implanted values does not undermine their moral responsibility. Still, because the manipulation makes the agent much more constitutively lucky, the agent’s degree of moral responsibility is seriously mitigated and can only be restored to pre-manipulation levels over time in the usual way—as described above—and so the issue of how moral responsibility is restored is built into the history-sensitive structuralist’s position.

The second argument for preferring history-sensitive structuralism is the argument from mere influence. Consider another of Mele’s cases, namely the case of Carl’s snacks:

Three months ago, Carl resolved not to eat any snacks for six months, mainly to address his weight problem but also to test his willpower. On average, Carl—who has not yet been manipulated—has a couple of medium-strength urges for a snack per day. On about one day in ten, he acts on such an urge; usually, he successfully resists these urges. Just for the fun of it, a manipulator flashes a subliminal message at Carl that he knows will give him a medium-strength urge for a snack. Carl succumbs to the urge...is he morally responsible for eating the snack? Provided that he is morally responsible for eating

snacks when the medium-strength urges at work are not produced by manipulation, a plausible answer is *yes*. (2019: 36)

I agree that this sort of manipulation, which I will refer to as *mere influence*, does not undermine the targeted agent's moral responsibility.

However, on Mele's view, mere influence does not undermine moral responsibility, but full-fledged manipulation (like in Beth's case) *does* undermine moral responsibility. Given that difference, Mele's view must be that there is a breaking point at which mere influence, which does not affect the targeted agent's moral responsibility, turns into manipulation, which completely undermines the targeted agent's moral responsibility. We could imagine a sorites-type series of cases, beginning with Carl's snacks and ending with a case where Carl is manipulated into eating a snack like Beth is manipulated into reviewing the manuscript, with intermediate cases in which the probability that the subliminal message will have its intended effect is varied. Perhaps Mele could develop his view in a way that allowed for degrees of moral responsibility to track degree of influence, but for all he has said about the cases on the ends of that series, it may be that there is a sharp cutoff from full moral responsibility to no moral responsibility.

On the history-sensitive structuralist position, by contrast, there is a ready explanation for the degrees of moral responsibility that would seem to track degree of influence/manipulation. It is open to the history-sensitive structuralist, at the very least, to say that this sort of influence, as it is ramped up, affects the agent's degree of moral responsibility in just the same way as manipulation does, but without reducing the agent to the seriously mitigated degree of moral responsibility had by entirely constitutively lucky agents. To the extent that subliminal influence introduces an element of luck, whether it be constitutive luck or some other form of luck, the

agent's degree of control is mitigated and for that reason their degree of moral responsibility is mitigated too.

5. History-Sensitive Structuralism vs. Standard Structuralism

At this point, my reasons for preferring history-sensitive structuralism over standard versions of structuralism may be obvious, but it is worth briefly making the contrast between the types of view explicit. On the standard structuralist position, manipulated agents are fully morally responsible for what they are manipulated to do. Going back to the case of Ann and Beth, these two philosophers are equally morally responsible for staying to review the manuscript, despite the fact that Beth would not have done this were it not for the manipulation. One standard structuralist, Manuel Vargas, considers a variation of the Beth case in which she is manipulated into having evil values for just one day, and he reports the following:

My intuition is that Beth is responsible. And that's my intuition even if the next day she is back to being good ol' Beth. I think that at least on the day she is evil, that she really is responsible for those evils...Regardless of how Beth came to be that way, it seems to me that she had the sort of control that is required for moral responsibility, and boy, did she act with ill will. (2013: 300-301)

In a similar vein, David Brink considers a case of what he calls "enabling" manipulation:

Imagine that Fig developed a form of psychosis, which rendered her insufficiently reasons-responsive. In this condition, she developed various paranoid delusions that Professor Plum was a threat to her, but let's assume that she was too incompetent to act on this false belief. Now, imagine that Fig is transformed by manipulation that bypasses her agency into someone who is reasons-responsive...In her reasons-responsive state, Fig kills Plum...If at the time of killing Plum Fig is fully reasons-responsive, then she is

responsible for killing him...She is no more excused than someone who had always been reasons-responsive but killed out of dislike or distrust. (2021: 105)

Whether it is manipulation that radically reverses an agent's values, or whether it is manipulation that enables an agent to satisfy structuralist conditions, these structuralists maintain that the agents in question are fully morally responsible for what they are manipulated to do—just as morally responsible as non-manipulated variants of these agents.

On my view, that is the wrong result. It is highly counterintuitive that these agents are fully morally responsible. This is exactly what motivates many historicists into endorsing a historical condition on moral responsibility in the first place. Even if intuitions should not be decisive in settling a debate like this—and I do not think they should be—intuitions are a guide, and it is a virtue of a view that it can account for our intuitions. In this case, Vargas and Brink take on a cost, as they are “biting the bullet” and failing to account for our intuitions.¹⁶

The history-sensitive structuralist, by contrast, can get the structuralist what they want—the result that manipulated agents can be responsible—but without saying that they are *just as responsible* as their non-manipulated counterparts. Furthermore, and once again, history-sensitive structuralism offers an *explanation* of why and how manipulated agents' moral responsibility is mitigated (in terms of their constitutive luck). For these reasons, the history-sensitive structuralist has advantages over the standard structuralist position.

¹⁶ In a recent paper, Craig Agule (2021) argues that our intuitions about such cases can be explained by the fact that both blame and sympathy are fitting—and those responses conflict. “The conflict between the two fitting reactions [blame and sympathy] provides a rich explanation for our conflicted experience with the cases, with no need to appeal to compromised responsibility at all” (2021: 161). While this may help the standard, history-*insensitive* structuralist, I can take on this explanation of our intuitions too, but also maintaining that manipulated agents' degree of moral responsibility is mitigated.

Acknowledgements

Thanks to audiences at the University of Delaware, especially Seth Shabo, Kate Rogers, Joel Pust, Hsin-wen Lee, and Daniel Koltonski, and at the Ohio Northern University Online Conference on Free Will and Luck, especially Bob Hartman, Paul Russell, Justin Capes, Yiling Zhou, and Gordon Cooper. Special thanks to Carolina Sartorio and Michael McKenna for comments on an earlier version of the paper. Work on this paper was supported by a Samford University Faculty Development Grant.

References

- Agule, Craig. 2021. "Being Sympathetic to Bad-History Wrongdoers," *Pacific Philosophical Quarterly* 102: 147-169.
- Arpaly, Nomy. 2003. *Unprincipled Virtue: An Inquiry into Moral Agency*. New York: Oxford University Press.
- Brink, David. 2021. *Fair Opportunity and Responsibility*. New York: Oxford University Press.
- Coates, Justin, and Swenson, Philip. 2013. "Reasons-Responsiveness and Degrees of Responsibility," *Philosophical Studies* 165: 629-645.
- Cyr, Taylor. 2019. "Moral Responsibility, Luck, and Compatibilism," *Erkenntnis* 84: 193-214.
- Cyr, Taylor. 2020. "Manipulation and Constitutive Luck," *Philosophical Studies* 177: 2381-2394.
- Cyr, Taylor. 2021. "The Inescapability of Moral Luck," *Thought* 10: 302–310.
- Feinberg, Joel. 1986. *Harm to Self*. Volume Three: *The Moral Limits of the Criminal Law*. New York: Oxford University Press.
- Fischer, John Martin. 2012. *Deep Control: Essays on Free Will and Value*. New York: Oxford University Press.
- Fischer, John Martin, and Ravizza, Mark. 1998. *Responsibility and Control: A Theory of Moral Responsibility*. Cambridge: Cambridge University Press.
- Frankfurt, Harry. 1988. *The Importance of What We Care About*. Cambridge University Press.
- Hartman, Robert. 2018. "Constitutive Moral Luck and Strawson's Argument for the Impossibility of Moral Responsibility," *Journal of the American Philosophical Association* 4: 165-183.

- Hartman, Robert. 2023. "Free Will and Moral Luck," in J. K. Campbell, K. Mickelson, and A. White (eds.), *The Blackwell Companion to Free Will* (Wiley Blackwell).
- Herdova, Marcela. 2020. "Barking Up the Wrong Tree: On Control, Transformative Experiences, and Turning Over a New Leaf," *The Monist* 103: 278-293.
- Levy, Neil. 2011. *Hard Luck: How Luck Undermines Free Will and Moral Responsibility*. New York: Oxford University Press.
- McKenna, Michael. 2004. "Responsibility and Globally Manipulated Agents," *Philosophical Topics* 32: 169-192.
- McKenna, Michael. 2012. "Moral Responsibility, Manipulation Arguments, and History: Assessing the Resilience of Nonhistorical Compatibilism," *Journal of Ethics* 16: 145-174.
- McKenna, Michael. 2021. "On Manipulated Agents and History-Sensitive Compatibilism," *Criminal Law and Philosophy* 15: 285-298.
- Mele, Alfred. 1995. *Autonomous Agents: From Self-Control to Autonomy*. New York: Oxford University Press.
- Mele, Alfred. 2006. *Free Will and Luck*. New York: Oxford University Press.
- Mele, Alfred. 2010. "Moral Responsibility for Actions: Epistemic and Freedom Conditions," *Philosophical Explorations* 13: 101–111.
- Mele, Alfred. 2019. *Manipulated Agents: A Window to Moral Responsibility*. New York: Oxford University Press.
- Mele, Alfred. 2020. "Moral Responsibility and Manipulation: On a Novel Argument Against Historicism," *Philosophical Studies* 177: 3143-3154.
- Mele, Alfred. 2021. "Manipulated Agents: Replies to Fischer, Haji, and McKenna," *Criminal Law and Philosophy* 15: 299–309.
- Mele, Alfred. Forthcoming. "Fischer on Epistemic and Freedom Requirements for Moral Responsibility," in T. Cyr, A. Law, and N. Tognazzini (eds.), *Freedom, Responsibility, and Value: Essays in Honor of John Martin Fischer* (Routledge).
- Nagel, Thomas. 1979. *Mortal Questions*. New York: Cambridge University Press.
- Nelkin, Dana Kay. 2016. "Difficulty and Degrees of Praiseworthiness and Blameworthiness," *Noûs* 50: 356-378.
- Paul, L. A. 2014. *Transformative Experience*. New York: Oxford University Press.
- Strawson, Galen. 1994. "The Impossibility of Moral Responsibility," *Philosophical Studies* 75: 5-24.

- Vargas, Manuel. 2013. *Building Better Beings: A Theory of Moral Responsibility*. New York: Oxford University Press.
- Watson, Gary. 1987. "Responsibility and the Limits of Evil: Variations on a Strawsonian Theme," in F. Schoeman (ed.), *Responsibility, Character and Emotions: New Essays on Moral Psychology* (Cambridge University Press).
- Watson, Gary. 2004. *Agency and Answerability: Selected Essays*. New York: Oxford University Press.