

Manipulation and Constitutive Luck

Taylor W. Cyr

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

Abstract: I argue that considerations pertaining to constitutive luck undermine historicism—the view that an agent’s history can determine whether or not she is morally responsible. The main way that historicists have motivated their view is by appealing to certain cases of manipulation. I argue, however, that since agents can be morally responsible for performing some actions from characters with respect to which they are entirely constitutively lucky, and since there is no relevant difference between these agents and agents who have been manipulated into acting from a character bestowed upon them by their manipulators, we should give up historicism. After presenting this argument and defending it against some potential objections, I briefly criticize the standard structuralist alternative and propose a new structuralist position that is shaped by reflection on constitutive luck.

Keywords: constitutive luck, historicism, manipulation, moral responsibility, structuralism

1. Introduction

According to several influential compatibilist accounts of moral responsibility, whether an agent is morally responsible for an action depends only on her psychological structure at that time (and not, say, on how she came to have that structure).¹ One of the main objections to such “structuralist” compatibilist accounts is that it is possible for the account’s allegedly sufficient conditions on moral responsibility to be satisfied by agents who have been manipulated—and who, thus, appear not to be morally responsible for their behavior. Partly to avoid this worry, “historicist” compatibilist accounts posit a historical condition on moral responsibility according to which an agent’s history (how she came to be a certain way) can make a difference as to whether or not she is morally responsible. Some authors proffer negative conditions that specify that morally responsible agents must *lack* a certain kind of history. Others defend positive conditions,

¹ See especially Frankfurt (1988) and Watson (2004). In this paper, I am concerned with direct (or non-derivative) moral responsibility. Since everyone should agree that there is a distinction between direct and indirect moral responsibility, and since this distinction is a historical one, everyone should think that, in cases of *indirect* moral responsibility, whether or not an agent is morally responsible can depend on non-structural features of an agent. For more on this point, see McKenna (2012: 156).

maintaining that morally responsible agents must *have* a certain kind of history. Proponents of both versions of historicism agree, though, that there is more to moral responsibility than merely satisfying certain structural conditions.

As I will argue in this paper, however, considerations pertaining to constitutive luck pose a challenge to historicism. Constitutive luck—a term introduced by Thomas Nagel (1979: 28)—is luck in being the kind of person that you are, or in having the character that you have. Briefly, the challenge to historicism is this: if, as historicists should admit, agents can be morally responsible for performing some actions (at the very least, the first actions for which they are morally responsible) from characters with respect to which they are entirely constitutively lucky, and since there is no relevant difference between these agents and agents who have been manipulated into acting from a character bestowed upon them by their manipulators (in particular, both are entirely constitutively lucky), then we should think (*contra* historicism) that such manipulated agents are just as morally responsible for their post-manipulation behavior as are non-manipulated but entirely constitutively lucky agents.²

² It is worth making two notes about this argument from constitutive luck. The first is that it shares a key structural feature with the increasingly popular “manipulation argument” against compatibilism (see especially Pereboom 2001, chapter 4 and 2014, chapter 4). Both use a “no relevant difference” premise, but whereas the manipulation argument aims to use this premise as a bridge from the non-responsibility of manipulated agents to the non-responsibility of ordinary determined agents, my “no relevant difference” premise aims to construct a bridge from the responsibility of constitutively lucky agents to the responsibility of (certain) manipulated agents.

Second, although I will continue to talk as though the historicist/structuralist debate is an in-house debate among compatibilists, the argument from constitutive luck has implications for *all* accounts of moral responsibility. Interestingly, although virtually all of the debate concerning whether moral responsibility is essentially historical has been a debate among compatibilists, even incompatibilist accounts of moral responsibility (i.e., libertarian accounts) may be divided into historicist and structuralist camps. (For more on this point, see McKenna 2016: 87-88. Kane’s 1996 libertarian account of freedom and responsibility is a nice example of a historicist libertarian account.) This is because it is possible for two agents to have identical psychological structures but different histories even if both agents inhabit indeterministic worlds. Given this, the argument

But reflection on constitutive luck should lead us not only to give up historicism but also to reject the extant version of structuralism on offer. According to the extant version of structuralism, manipulated agents are just as responsible for what they do as are relevantly similar non-manipulated agents, but manipulated agents are not responsible for having the character from which they act, whereas non-manipulated agents may well be. In other words, on this view, even though manipulated agents are just as responsible *for what they do* as are non-manipulated agents, non-manipulated agents are responsible *for more things* than manipulated agents. As I argue in this paper, however, such an account fails to take seriously the extent to which constitutive luck can *mitigate* moral responsibility. Given this problem for the extant version of structuralism, I will propose a new structuralist position that is shaped by reflection on constitutive luck and can account for the mitigating effects of such luck on moral responsibility.

I will proceed as follows. In section 2, I will discuss a case of manipulation as well as some examples of historical conditions on moral responsibility. I will then, in section 3, develop the argument from constitutive luck in more detail, and, in section 4, I will consider some objections to the argument and offer replies to them. Next, in section 5, I will argue that the standard version of structuralism is, like historicism, undermined by considerations pertaining to constitutive luck, and, finally, in section 6, I will introduce a new structuralist position—a history-sensitive structuralism—that provides a new account of the responsibility of certain manipulated agents.

2. Manipulation and Historicism

To see how manipulation cases typically motivate compatibilists to be historicists, consider the following story and commentary from Alfred Mele:

from constitutive luck should persuade not only *compatibilists* but also *libertarians* to plump for structuralist accounts of moral responsibility.

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 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...

Ann, by hypothesis, freely does her philosophical work, but what about Beth? ...the difference in their current status regarding freedom would seem to lie in how they *came* to have certain of their psychological features, hence in something *external* to their present psychological constitutions. That is, the crucial difference is *historical*; free agency is in some way history-bound. (2006: 164-166)³

³ This example is based on the one in Mele (1995: 145). Whereas I have been (and will continue to be) focused on *moral responsibility*, Mele typically focuses on *freedom*. Nevertheless, we are talking about the same features of agency, for, as Mele notes concerning his use of the term “free action” and its cognates:

My interest is in what might be termed *moral-responsibility-level free action*—roughly, free action of such a kind that if all the freedom-independent conditions for moral responsibility for a particular action were satisfied without that sufficing for the agent's being morally responsible for it, the addition of the action's being free to this set of conditions would entail that he is morally responsible for it. (2006: 17)

Suppose that, 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. By stipulation, Beth satisfies all structuralist conditions on moral responsibility when she does this, but many find it intuitive that Beth is not morally responsible for this action. If Beth is not morally responsible despite being the time-slice duplicate (or “psychological twin”) of a typical morally responsible agent, then some version of historicism is true.

According to Mele’s own historicist proposal, morally responsible agents must *lack* histories of a certain kind, which is to say that Mele is a *negative* historicist. More specifically, here is Mele’s historical condition on free actions (and thus, he thinks, on moral responsibility), which is built into (the moderate version of) his compatibilist proposal:

An agent [performs some action *A*] freely if he nondeviantly *A*-s on the basis of a rationally formed deliberative judgment that it would be best to *A*, has no compelled or coercively produced attitudes that influence his deliberative judgment, is well informed on the topic of his deliberation, and is mentally healthy. (2006: 200)⁴

On this proposal, an agent who is morally responsible for *A*-ing must lack a history in which compelled or coercively produced attitudes influenced her deliberative judgment (about *A*-ing). In the case considered above, when Beth forms the deliberative judgment that it would be best to stay at the office to review a manuscript for a journal, this judgment is influenced by attitudes that she

⁴ Given its requirement that the agent have rationally formed a deliberative judgment, it may seem as though Mele’s proposal is really a positive (or a “mixed”) historicist account, requiring that morally responsible agents *have* a certain sort of history (and, on a mixed view, also *lack* a certain sort of history). But the requirement that morally responsible agents have rationally formed a deliberative judgment is one that can be satisfied at (or just before) the time of action and so, arguably, is part of the (perhaps slightly temporally extended) time of action that structuralists take moral responsibility to depend on. Thanks to John Fischer for encouraging me to address this point.

has been compelled to have. On Mele's view, then, Beth is not morally responsible for her post-manipulation behavior.

Unlike Mele, John Martin Fischer and Mark Ravizza (1998) defend a *positive* historicist view. More specifically, on their view, in order for an agent to be morally responsible for an action, the agent must have *taken responsibility* for the mechanism (e.g., practical reasoning) that issued in the action (1998: 210-217). Fischer and Ravizza claim that there are three (necessary and sufficient) conditions on taking responsibility:

First, the agent must view himself – when acting from certain mechanisms – as an agent; he must see that certain upshots in the world are the results of his choices and actions. Second, an agent must view himself as an apt target for the reactive attitudes... Finally, the cluster of beliefs specified by the first two conditions must be based, in an appropriate way, on the individual's evidence. (1998: 238)

On this proposal, an agent who is morally responsible for *A*-ing must have a history in which she took responsibility for the mechanism from which she acts in *A*-ing. Now, according to Fischer and Ravizza, agents who have been manipulated act from a "manipulation mechanism" (1998: 230-236), which is distinct from (say) ordinary practical reasoning, for which they have presumably taken responsibility. Even if a manipulated agent satisfies the first two conditions on taking responsibility, the agent would not, Fischer and Ravizza claim, satisfy the third—her beliefs about herself as an agent that is an apt target for the reactive attitudes would not be *appropriately* based on her evidence, given the manipulation (1998: 235-236). So, on their view, because Beth's staying late to review the manuscript issues from a mechanism for which she has not taken responsibility, Beth is not morally responsible for her post-manipulation behavior.

3. Constitutive Luck

We have seen a typical case of manipulation that serves to motivate historicism, and we have seen a version of a negative historical condition as well as a version of a positive historical condition on moral responsibility. In this section, I will present an argument against all versions of historicism. The argument has two premises, one claiming that agents who are entirely constitutively lucky can be morally responsible for what they do, and the other claiming that manipulated agents like the ones we have been considering are not relevantly different from agents who are entirely constitutively lucky. What follows is that manipulated agents like the ones we have been considering can be morally responsible for what they do, despite failing to satisfy historical conditions (whether positive or negative) on moral responsibility, which is to say that historicism is false. Here, then, is the argument from constitutive luck:

Premise 1: Agents who are entirely constitutively lucky can be morally responsible for what they do.

Premise 2: There is no relevant difference between agents who have been manipulated in certain ways and agents who are morally responsible for actions that stem from characters with respect to which they are entirely constitutively lucky.

Conclusion: Manipulated agents like the ones we have been considering can be morally responsible for what they do, despite failing to satisfy historical conditions (whether positive or negative) on moral responsibility, which is to say that historicism is false.

If compatibilism is true (rather than, say, skepticism about moral responsibility) and this argument is sound, then structuralism is true. After defending each of the premises of the argument from constitutive luck in the remainder of this section, I will, in the next section, consider some

objections to the argument and will offer replies to them. Even if my replies to these objections succeed, however, the manipulation objection to structuralism remains. In the remainder of the paper, then, after a brief discussion of the extant structuralist response to the manipulation objection, I propose a new structuralist position that handles the objection differently. In addition to being better suited to handle the manipulation objection to structuralism, the new position avoids a certain problem—one that stems from reflection on constitutive luck—for the standard structuralist view.⁵

Given that we are finite agents (and, in particular, given that we began to exist), if we are morally responsible for any of our actions, there was for each of us a first action for which we were morally responsible (hereafter “first morally responsible action”).⁶ Now, with respect to an agent’s first morally responsible action, she is entirely constitutively lucky in having the character from which that action stems. For an agent to be entirely constitutively lucky minimally requires that she not have exercised control over the formation of her character.⁷ Of course, we do sometimes exercise control over the formation of our characters, as when, in another of Mele’s

⁵ As I noted in section 1, this argument from constitutive luck is structurally similar to the increasingly popular manipulation argument against compatibilism. But notice that, despite sharing the structural similarity of having a “no relevant difference” premise to bridge the first premise to the conclusion, it is not the case that proponents of one of these arguments must endorse the other. In fact, while I defend the argument from constitutive luck that I develop in this paper, I have argued elsewhere that compatibilists should reject the first premise of the manipulation argument against compatibilism. For my take on that argument, see Cyr (Manuscript).

⁶ One might think that this boundary at the beginning of our careers as agents (between actions for which we are not morally responsible and actions for which we are morally responsible) is a vague one. (We can all agree, of course, that agents *gradually* come to have the capacities necessary for moral responsibility; what we might disagree about is whether there is a precise threshold at which one performs one’s *first* morally responsible action.) While this would suffice to undermine my argument as I present it in the main text, the argument could, albeit with some clunkiness, be modified to refer to the first action outside of that zone of vagueness.

⁷ This is only a necessary condition on being entirely constitutively lucky, and this much should be uncontroversial.

cases, Chuck resolves to become a person “who unambivalently rejects conventional morality as a system designed for and by weaklings” (2006: 171) and intentionally brings about this aspect of his character by performing a series of cruel actions that have heart-hardening effects. After this process, Chuck is not entirely constitutively lucky, for he has exercised some control over the formation of his character. But no agent exercises such control over the formation of her character *before* performing her first morally responsible action, or else the action in question wouldn’t be her *first* morally responsible action.⁸ If our career as morally responsible agents is to get off the ground in the first place, it must be the case that we can be morally responsible for actions that stem from characters with respect to which we are entirely constitutively lucky. Thus, anyone who accepts (as historicists must) that moral responsibility gets off the ground should also accept that we can be morally responsible while entirely constitutively lucky, and this would be to accept the first premise of the argument from constitutive luck.⁹

But, as it turns out, 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. Recall the circumstances surrounding Beth’s post-manipulation behavior. When Beth stays late to review the manuscript, she does so because of her new hierarchy of values, aspects of her character over which she exercised no control, and thus her action stems from a character with respect to which

⁸ Might we exercise control over the formation of our characters, not be morally responsible for the changes to our characters, and yet not be entirely constitutively lucky? I do not think this is possible, for 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.

⁹ As an anonymous reviewer points out, the argument of this paragraph parallels an argument from Pereboom (2001: 48-49) against Kane’s (1996) historical libertarian account.

she is entirely constitutively lucky. Yet anything that might be said in favor of attributing moral responsibility to an agent at the time of her first morally responsible action will apply to Beth as well: she is mentally healthy, rationally competent, aware of the moral significance of the action, etc. And notice that anything that might count against Beth's being morally responsible for her post-manipulation behavior should apply just as well to agents at the time of their first morally responsible actions: they have been "given" characters over which they had no say, they are not morally responsible for endorsing their characters, etc. These considerations provide a good amount of support for the second premise of the argument from constitutive luck.

Yet if agents can be morally responsible while entirely constitutively lucky (as the first premise maintains), then, given that there is no relevant difference between such agents and agents who are manipulated like Beth is (as the second premise maintains), manipulated agents like Beth can be morally responsible for their post-manipulation behavior, too. But since agents like Beth do not satisfy (either positive or negative) historical conditions on moral responsibility, historicism is false. Since historicists claim that we (finite agents) are morally responsible for some of what we do, historicists who are not convinced by the argument from constitutive luck will have to object to its second premise. In the next section, I will address three potential objections to this premise.

4. Objections and Replies

According to one potential objection to the argument from constitutive luck, the relevant difference (between Beth and an agent at the time of her first morally responsible action) is a difference in the seriousness (or moral severity) of their respective actions. While Beth is an adult agent, the objector might continue, our first morally responsible actions are those that we performed when we were much younger—actions like insulting a classmate or disobeying our

parents—and that were much less serious than the actions that we (and the manipulated agents in the cases we discuss) perform as adults. Since there is a relevant difference, the objection concludes, there is room to make asymmetric attributions of moral responsibility about younger, “little agents,” on the one hand, and manipulated agents, on the other.¹⁰

While this potential objection is a natural response, there are two reasons that I do not think it can succeed. First, it is not obvious to me that there is a difference in the seriousness of the respective actions. Is insulting a classmate less serious than staying late to review a manuscript? If anything, the former seems more serious to me than the latter, especially taking into account the possible long-term, character-forming effects of the actions we perform at a younger age. But, setting this aside, a second worry for this objection is that we could alter the case of manipulation, replacing the manipulated action with one that is typical for younger agents to perform (such as a lie to a parent, or a demonstration of kindness to a sibling), with no change to the conclusion of the argument. Imagine that, after her manipulation, rather than staying late to review a manuscript, Beth tells a (minor) lie to her father. Such an action does not differ in seriousness from candidates for young agents’ first morally responsible actions. So, if young agents can be morally responsible for actions with exactly the same degree of seriousness, as compatibilists should admit, manipulated agents may be morally responsible for their post-manipulation behavior (despite failing to satisfy proposed historical conditions on moral responsibility). Thus, I do not think that a difference in the seriousness of the actions will suffice to block the no-relevant-difference premise of the argument.

A second objection maintains that the relevant difference (between Beth and an agent at the time of her first morally responsible action) is a difference in the capacities of the agents. In

¹⁰ Thanks to Al Mele for suggesting this objection.

particular, according to this objection, while Beth has been manipulated into acting from a character over which she had no control, she retains her ability to recognize and respond to reasons, her competency for social interaction, her impulse-control, etc. Little agents, by contrast, lack these capacities or, at the very least, do not possess them to the same degree as manipulated agents like Beth. Given this difference, the objection concludes, we should deny the “no relevant difference” premise of the argument from constitutive luck.¹¹

I have two things to say in reply to this objection. First, if the difference in the capacities of the two types of agents makes a difference here, presumably the type of agent who *lacks* the capacities (or who possesses them to a lesser degree) is the candidate for non-responsibility, not the agent who *has* the capacities. Yet, according to the objection, the *manipulated agents* are the ones who possess the greater capacities. Even if there is a difference to these agents’ degree of moral responsibility, then, the compatibilist, insofar as she accepts that the agents with lesser capacities can be morally responsible, should admit that manipulated agents (who fail to satisfy historical conditions on moral responsibility) can be morally responsible. In other words, even if we grant the point that the difference in capacities is a moral responsibility-relevant difference between the two types of agents, the argument from constitutive luck could run with the following in place of the original version of the second premise (the “no relevant difference” premise):

*Premise 2**: If agents who are entirely constitutively lucky can be morally responsible for what they do, then agents who have been manipulated in certain ways can be morally responsible for actions that stem from characters that they were manipulated into having.

¹¹ Thanks to Ben Matheson for suggesting this objection.

Since the compatibilist must accept that agents who are entirely constitutively lucky can be morally responsible for what they do, pointing out that such agents have lesser capacities than manipulated agents like Beth will not help to escape the conclusion of the argument from constitutive luck.

Second, *even if* the difference in the capacities of the two types of agents originally considered *is* a relevant difference, this would not undermine the argument from constitutive luck. To see why not, consider a new pair of cases. In the first case, “little Tony” performs his first morally responsible action (and so he is entirely constitutively lucky but has the capacities required to be morally responsible) when he shows kindness to his sister.¹² In the second case, imagine that little Anton, a time-slice duplicate of little Tony, performs the same action but, as it turns out, had just previously been manipulated to have the same hierarchy of values as Tony.¹³ By stipulation, both Tony and Anton have exactly the same capacities, and yet Anton fails to satisfy historical conditions on moral responsibility. Another way to put the objection to historicism that is codified in the argument from constitutive luck, then, is to ask: on what basis could the historicist maintain that little Tony can be morally responsible for his action but little Anton cannot be morally responsible for his? Since historicists must admit that agents like little Tony can be morally responsible for their actions, and since there is no relevant difference between Tony and Anton, historicists should accept that Anton can be morally responsible for his action despite failing to satisfy historical conditions on moral responsibility, which is to say that they should give up their historicism.

Finally, a third objection says that the relevant difference (again, between Beth and an agent at the time of her first morally responsible action) is that manipulated agents presumably had

¹² This example is based on one from Mele (2006: 129-130).

¹³ For a related but slightly different use of this example, see section 4 of Cyr (2019).

a say over the way their characters *were* (before the manipulation), whereas agents who are entirely constitutively lucky at the *start* of their careers as morally responsible agents have not *lost* some character over which they previously had some say.

I have two replies to this objection, too. First, if this objection were to succeed (i.e., if this really is a *relevant* difference between the agents), then it would prove too much, so to speak. To see why, consider some of the examples of “conversions” introduced by Nomy Arpaly:

There are cases of profound conversions in a person's values that are just as inexplicable to their possessors, just as uninvited, and just as irrational as Beth's conversion. For instance, many people have undergone shifts from being self-endorsed party animals to being self-endorsed industrious workers because of mysterious factors they regard as “age” or the “drying up of hormones.” Other people begin to value parenthood—value, not just like—the moment their (formerly unwanted) children are born. Other people have converted from atheism to religion (or vice versa) as the result of an experience of extreme loneliness and pain, and many people's convictions, however well rationalized, are the result of irrational or a-rational factors so unconscious that the agent had about as much opportunity to subject them to reflective criticism as Beth had to scrutinize her brainwashing process. (2003: 127)

As these examples show, there are many real-life cases in which a person comes to repudiate aspects of her character that she had exercised control in performing, yet we do not think that such agents could not be morally responsible for acting from new values simply because those values stand in opposition to the values the agent previously had.¹⁴ So it is possible for an agent to have

¹⁴ Of course, as Arpaly (2003: 128) explains, unlike the agents in these other examples, Beth's autonomy is violated. But, as Arpaly (2003: 129) argues, it does not follow that, unlike the others, Beth is not morally responsible for what she does. See also Fischer (2012). Indeed, I would endorse

some say over her character and then to be subjected to a modification over which she had no say without thereby losing moral responsibility for her subsequent actions.

Second, even if we are inclined to think that being subjected to some modification to our characters (whether by another agent or by no agent at all) would mitigate our moral responsibility for what we did upon having the new character, presumably we would nevertheless admit that, at some point after the change to character, we could eventually perform some morally responsible action. Indeed, some historicists have made exactly this claim.¹⁵ But if we can eventually perform some morally responsible action, what would need to happen in the interim in order for this to be possible? I submit that it would be odd to require that, after such a change to character, an agent must satisfy conditions that are not satisfied by agents who perform their first morally responsible actions. Instead, compatibilists should admit that agents like Beth can be morally responsible right after being subjected to a modification in character (though, as we will see in the remainder of the paper, they should also say that agents like Beth act with mitigated moral responsibility upon being manipulated).

I have replied to three objections to the second premise of the argument from constitutive luck. A different response to the argument is to deny its first premise, which in this context (and given that we—finite agents—began to exist) would be tantamount to endorsing moral responsibility skepticism. Insofar as the historicist was already committed to compatibilism, however, endorsing the argument's conclusion and opting instead for a structuralist position would require less revision to her theory, so my advice to historicists is to convert to structuralism. Still,

Arpaly's charge: "Anyone who wishes to argue that Beth is not morally responsible for her actions would need to explain why having been irrationally influenced by an evil human being exempts from responsibility in a way that having been influenced in a similar way by some unlucky chance of a force of nature does not" (2003: 129).

¹⁵ For example, see Fischer and Ravizza (1998: 235).

as we saw above, structuralism is vulnerable to the manipulation objection. How should structuralists handle this objection?

5. Standard Structuralism and the Manipulation Objection

One way for structuralists to respond to the challenge from manipulation cases like Mele's Ann/Beth case is to point out that while Beth is morally responsible for doing what she is manipulated to do, she is not morally responsible for having the character she was manipulated to have. When Ann acts, by contrast, she is morally responsible for more things—both for having the moral personality from which her action ensues and also for her action itself. As Michael McKenna (2012: 163) puts the point on behalf of structuralism (though he himself remains agnostic with respect to this debate), perhaps we can account for the “intuitive unease” of taking Beth to be just as morally responsible for what she does as Ann is for what she does by noting that, though both are morally responsible for what they do, Ann is morally responsible for more than Beth is. This difference can be highlighted by considering McKenna's “instant agent” thought experiment. Suppose that Suzie has just been created “to be a psychologically healthy woman indistinguishable from any other normally functioning thirty-year-old person” (2004: 180), and suppose further that she is an intrinsic duplicate of Ann, complete with false beliefs about a past filled with philosophical activity. When Suzie stays late to review a manuscript (just after her creation), even if she is morally responsible for this action, clearly she is not morally responsible for coming to have the character from which it stems. Thus, according to McKenna's proposal, Suzie is just as morally responsible for staying late to review the manuscript as Ann, yet Ann is morally responsible for more than just her action of staying late to review the manuscript.

One problem for this proposal, however, is that it does not take seriously enough the way that constitutive luck mitigates moral responsibility.¹⁶ Consider another pair of “little agents.” Suppose again that little Tony is entirely constitutively lucky and yet morally responsible for showing kindness to a sibling. Now imagine another agent, little Anthony, who performs the same action as Tony and who is exactly like Tony except that Anthony is not entirely constitutively lucky—perhaps he has previously endorsed his value of sibling-friendship. Insofar as constitutive luck mitigates control, and insofar as control is necessary for moral responsibility, little Tony would appear to be less morally responsible than little Anthony *for the action of showing kindness to a sibling* (not just morally responsible for fewer things than Anthony) despite the fact that the only difference between Tony and Anthony is a difference in the extent to which they are constitutively lucky. For this reason, though I think Suzie is just as morally responsible for staying late to review the manuscript as is Beth (as both are constitutively lucky to the same degree), I do not think that either of these agents is *as* morally responsible for this action as Ann is, given that they are entirely constitutively lucky with respect the characters that lead to that action. And the difference between them and Ann does not reduce to a difference in how many things each is morally responsible for; rather, because of their constitutive luck, Beth and Suzie are more like little agents performing their first morally responsible actions. So, if we take seriously the extent to which we can be constitutively lucky (and if we do not take this luck to undermine our moral responsibility for what we do), we should allow that constitutive luck can mitigate our moral responsibility.

¹⁶ Here and in my own proposal I presuppose that *moral responsibility* comes in degrees. Some deny this presupposition, maintaining instead that *praiseworthiness and blameworthiness* come in degrees. What I say here could be adopted, *mutatis mutandis*, by someone who held this alternative view. Thanks to John Fischer for suggesting that I mention this variant. For a further elaboration of my view, which I will only sketch here, see Cyr (2019).

6. A New Proposal: History-Sensitive Structuralism

This problem for the extant version of structuralism suggests an alternative proposal: an agent's history cannot (by itself) undermine her moral responsibility for what she does, but it can affect the *degree to which* she is morally responsible for what she does. This proposal is a structuralist one, since it maintains that structural features of an agent are sufficient to establish that she meets the threshold conditions on moral responsibility (*contra* historicism), but, unlike the extant structuralist position, the proposal nevertheless allows for an agent's history to affect (the degree of) her moral responsibility for what she does (and not just how many things she is morally responsible for). It is, thus, a *history-sensitive structuralism*. According to this proposal, since Tony satisfies the threshold conditions on moral responsibility, he is (at least a little bit) morally responsible, but Anthony may be *more* morally responsible for the same action because of his history—in particular, a history in which he has endorsed a pertinent value, thereby mitigating his constitutive luck at the time of the action. And in the case of manipulation considered above, since Beth satisfies typical structuralist conditions on moral responsibility even after she is manipulated, this history-sensitive structuralist proposal says that we should grant that Beth is morally responsible for staying to review the manuscript; given that Beth has been manipulated, however, and in particular because of the fact that Beth is constitutively lucky in having the character from which this action stems, the present proposal says that Beth is less morally responsible for staying to review the manuscript than Ann is.

This is, of course, only a sketch of an alternative structuralist position. One lingering question for the sketch concerns exactly how one's history interacts with the standard conditions on moral responsibility (the epistemic and control conditions) in determining an agent's degree of moral responsibility. It is open to those who accept history-sensitive structuralism to part ways in

their further developments of the position. The aim of this paper, though, has been to show that considerations pertaining to constitutive luck give us reason to reject all versions of historicism as well as the standard form of structuralism, which leaves history-sensitive structuralism as our only non-skeptical option.

7. Conclusion: Counting the Cost

I have argued that reflection on constitutive luck should lead historicists to convert to structuralism. As we have seen, though, a major challenge to structuralism is the manipulation objection: manipulated agents, even if they satisfy structuralist conditions on moral responsibility, are intuitively not morally responsible for their post-manipulation behavior. Some structuralists have “bitten the bullet” and admitted the counterintuitive consequence of their view. For example, in a passage that is frequently cited in the debate between historicists and structuralists, Harry Frankfurt says:

...to the extent that a person identifies himself with the springs of his actions, he takes responsibility for those actions and acquires moral responsibility for them; moreover, the questions of how the actions and his identifications with their springs are caused are irrelevant to the questions of whether he performs his actions freely and is morally responsible for performing them. (1988: 54)

Given the manipulation challenge, however, a natural and common complaint against structuralism is that it “comes with a considerable cost” (McKenna 2012: 169).

On the history-sensitive structuralist account that I have introduced here, however, because manipulated agents who satisfy typical structuralist conditions are entirely constitutively lucky, they are only morally responsible for what they do *to a slight degree*, just as agents who perform their first morally responsible actions are only morally responsible to a slight degree. Insofar as

we are willing to grant that “little agents” (who are entirely constitutively lucky) can be *a little bit* morally responsible for what they do, it would seem less counterintuitive to maintain that manipulated agents can likewise be a little bit morally responsible than it would be to maintain that they are *just as* responsible for what they do as if they had not been manipulated. If history-sensitive structuralism comes with a cost, I conclude, it is not as high a cost as comes with rival structuralist proposals.

Acknowledgements

For comments on earlier drafts of this paper, I am grateful to John Fischer and Michael Nelson; the members of the Agency Workshop at the University of California, Riverside, especially Zac Bachman, Dave Beglin, Andrew Law, Meredith McFadden, Jonah Nagashima, and Debbie Nelson; my commentator at the 2017 Pacific APA in Seattle, Ben Matheson; the audience at that APA, especially Craig Agule, Garrett Pendergraft, Michael Robinson, Philip Swenson, and Neal Tognazzini; and to an anonymous reviewer for this journal. I am also grateful to Al Mele for discussion of the main argument of this paper.

References

- Arpaly, Nomy. (2003) *Unprincipled Virtue: An Inquiry into Moral Agency*. New York: Oxford University Press.
- Cyr, Taylor. (2019) “Moral Responsibility, Luck, and Compatibilism,” *Erkenntnis* 84: 193-214.
- Cyr, Taylor. (Manuscript) “Manipulation Arguments and Libertarian Accounts of Free Will.”
- Fischer, John Martin. (2012) “Responsibility and Autonomy: The Problem of Mission Creep,” *Philosophical Issues* 22: 165-184.
- 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.
- Kane, Robert. (1996) *The Significance of Free Will*. 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. (2016) “A Modest Historical Theory of Moral Responsibility,” *Journal of Ethics* 20: 83-105.
- 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.
- Nagel, Thomas. (1979) “Moral Luck,” in *Mortal Questions*. Cambridge: Cambridge University Press.
- Pereboom, Derk. (2001) *Living Without Free Will*. Cambridge: Cambridge University Press.

- Pereboom, Derk. (2014) *Free Will, Agency, and Meaning in Life*. New York: Oxford University Press.
- Watson, Gary. (2004) *Agency and Answerability: Selected Essays*. New York: Oxford University Press.