

5-2011

In Defense of the Four-Case Manipulation Argument for Hard Incompatibilism

Krasimira Filcheva

The College at Brockport, filcheva@live.unc.edu

Follow this and additional works at: <http://digitalcommons.brockport.edu/honors>

 Part of the [Philosophy Commons](#)

Repository Citation

Filcheva, Krasimira, "In Defense of the Four-Case Manipulation Argument for Hard Incompatibilism" (2011). *Senior Honors Theses*.
66.

<http://digitalcommons.brockport.edu/honors/66>

This Honors Thesis is brought to you for free and open access by the Master's Theses and Honors Projects at Digital Commons @Brockport. It has been accepted for inclusion in Senior Honors Theses by an authorized administrator of Digital Commons @Brockport. For more information, please contact kmyers@brockport.edu.

A Senior Honors Thesis

Presented in Partial Fulfillment of the Requirements
for graduation in the College Honors Program

By

Krasimira Filcheva
Philosophy Major

The College at Brockport
May 2011

Thesis Director: Gordon Barnes, Assistant Professor, Philosophy

Abstract

Derk Pereboom's four-case manipulation argument has proved to be a major point of contention between compatibilism and hard incompatibilism in the debate over causal determinism's alleged threat to free will and moral responsibility. Notably, the four-case argument has met Michael Mckenna's so called hard-line reply, a six-case argument modeled after Pereboom's four-case one and intended to establish a dialectical stalemate between the compatibilist and incompatibilist positions on largely intuitive grounds. Mckenna contends that his six-case argument elicits compatibilistically friendly intuitions about Pereboom's case 1 in which the agent is said to be morally responsible. I argue that Mckenna's hard-line reply does not succeed in demonstrating this alleged stalemate between the two debating sides. The current state of the dispute can be so characterized only if the opposing sides' reported intuitions enjoy evidential equivalence. But, I argue, the evidential credentials of these intuitions are not equivalent. The newly elicited intuition from Mckenna's six-case argument cannot be assumed to do any evidential work without a good explanation of why it cannot be a commanding intuition. I argue that Mckenna's proposed explanation is not adequate as it stands. Finally, I offer a diagnosis of the origins of the apparent dialectical impasse, and, on this basis, seek to advance the debate on novel grounds. I defend the four-case argument by locating the disagreement between the two sides in the particular application of the generalization method.

Table of Contents

Abstract.....	2
Acknowledgments.....	4
Chapter I.....	5
Introduction.....	5
The Four-case Argument.....	6
Michael Mckenna’s Hard-line Reply.....	8
Chapter II.....	11
Mckenna’s Explanation: An Appraisal.....	11
Chapter III.....	18
Dialectical Burdens.....	18
Chapter IV.....	25
A Diagnosis.....	25
Chapter V.....	29
Advancing the Debate.....	29
Bibliography.....	39

Acknowledgments

I would like to thank Professor Gordon Barnes for his invaluable support and constructive criticism in relation to this project. His continuing guidance and advice in the course of writing this paper has immeasurably enriched the quality of the work. I also remain grateful for Professor Donna Kowal's help and support in the beginning stages of the Honors thesis experience. I benefitted immensely from the sense of direction that she instilled in the members of the Junior Colloquium class, as I began thinking about the thesis project.

In Defense of the Four-Case Manipulation Argument for Hard Incompatibilism

Krasimira Filcheva

I. Introduction

Derk Pereboom's four-case manipulation argument has proved to be a major point of contention between compatibilism and hard incompatibilism in the debate over causal determinism's alleged threat to free will and moral responsibility. Notably, the four-case argument has met Michael Mckenna's so called hard-line reply, a six-case argument modeled after Pereboom's four-case one and intended to establish a dialectical stalemate between the compatibilist and incompatibilist positions on largely intuitive grounds. Mckenna contends that his six-case argument elicits compatibilistically friendly intuitions about Pereboom's case 1 in which the agent is said to be morally responsible. He thus seeks to demonstrate that the incompatibilist enjoys no theoretical advantage over the compatibilist. According to Mckenna, the debate seems to have reached a dialectical impasse, given that both positions report equally legitimate but conflicting intuitions about Pereboom's case 1.

In what follows, I argue that Mckenna's hard-line reply does not succeed in demonstrating this alleged stalemate between the two debating sides. The current state of the dispute can be so characterized only if the opposing sides' reported intuitions enjoy evidential equivalence. But, I argue, the evidential credentials of these intuitions are not equivalent. The newly elicited intuition from Mckenna's six-case argument cannot be assumed to do any evidential work without a good explanation of why it cannot be a commanding intuition. I argue that Mckenna's proposed explanation is not adequate as it stands. In the second part of the paper, I offer a suggestion on how to advance the debate over the four-case argument on novel grounds.

The Four-Case Manipulation Argument

Pereboom's argument attempts to show that there is no principled moral distinction between cases in which agents are causally determined to perform a given action and cases in which they are manipulated into performing the same action in a responsibility-undermining way. The argument seeks to elicit the intuition that causally determined agents are not morally responsible by way of a carefully crafted generalization strategy. It comprises a series of three manipulation cases in which an agent performs a morally objectionable action and ends with a fourth case in which regular causal processes lead to the performance of the same action. Here, the intuitive response to the first case, that the agent is not morally responsible, is supposed to generalize to the second one. Pereboom has structured the cases in such a way that there are, arguably, no morally relevant differences between them. The generalization strategy is then used for the second and third cases and, importantly, the third and fourth one, an uncontroversially natural case in which the agent's action is determined by the past and the laws of nature (Pereboom 93-8).

It is important to stress that, in constructing his four cases, Pereboom has incorporated all prominent compatibilist conditions that are usually taken to be sufficient for an action to be free – reasons-responsiveness, receptivity to moral reasons, the presence of second-order desires, etc. In other words, what Michael Mckenna refers to as the Compatibilist-friendly Agential Structure, CAS (142), is supposed to be as prominent as the relevant causal factors – manipulation and natural causal processes – in all four cases. The following, then, is a brief presentation of the cases.

In Pereboom's case 1, the manipulation is local. Sophisticated neuroscientists manipulate Plum from moment to moment by radio-technology so that he is causally determined to kill Ms. White, but in such a way that the compatibilist conditions for moral responsibility are met. Plum acts on the basis of his non-alienating desires and stable character without being constrained to act. His action thus fits the conditions for responsibility proposed by Hume (Hume 1739/1978, 319-412; Pereboom 94). His second-order desires are such that he wills his egoistic first-order desire to kill Ms. White and wants to act on the basis of it, thereby satisfying Frankfurt's condition for agent responsibility (Frankfurt 1971; Pereboom 94). Plum is also reasons-responsive in that he would be willing to refrain from killing Ms. White were he to regard the consequences as too costly (Fischer and Ravizza 69-82; Pereboom 94). His patterns of moral deliberation also conform to the responsibility conditions advanced by Jay Wallace. Plum is capable of guiding his behavior on the basis of moral reasons, which is evidenced by his tendency to give such reasons considerable weight when egoistic motives are not strong enough. He is also capable of subjecting his behavior to revision over time (Wallace 51-83; Pereboom 94).

In a second scenario, Plum is programmed by the neuroscientists from the beginning of his life so that his character develops in accordance with egoistic motivations that culminate in his killing Ms. White. The compatibilist agential structure necessary for moral responsibility is preserved as Plum appears to be just like any ordinary agent. In the third scenario, Plum is determined to kill Ms. White by the rigorous psychological conditioning of his community over which he has no control. But, the hard incompatibilist tells us, the difference between this and the previous scenario is just that the manipulation is more ordinary in this case. That is hardly a principled ground on which to distinguish between the cases and hold Plum responsible in the

third one. Non-responsibility thus generalizes. In Case 4, Plum is an ordinary agent in our natural world, causally determined to kill White by the past and the laws of nature but in such a way that his action conforms to the compatibilist conditions for responsibility. Since there seems to be no good explanation of why Plum may be thought responsible in the last scenario but not in the third, the incompatibilist concludes that this is just to say that Plum is not morally responsible for his action even when it is produced by normal causally deterministic processes (Pereboom 94-8).

Here, the judgment that manipulated Plum is not morally responsible in case 1, based on a widely shared and strong intuition, is conjoined with Pereboom's generalization strategy mentioned above to elicit the same judgment in the ordinary case 4. The success of the incompatibilist argument thus depends in large measure on the original intuition that manipulated Plum in case 1 is clearly not morally responsible. If the compatibilist can elicit a different intuition to counteract the one reported by Pereboom, then the argumentative force of the four-case argument can seem dispelled. This line of reply is precisely the one employed by Mckenna.

Michael Mckenna's "Hard-Line" Reply¹

Proceeding in a reverse way, from a sixth case involving regular deterministic processes to Pereboom's case 1 and applying Pereboom's own generalization strategy, Mckenna argues in the following way. Since the incompatibilist cannot assume that determinism undermines moral responsibility, he or she cannot deny that the agent is free and morally responsible in Mckenna's normal case 6 (modeled after Pereboom's case 4). Here, Mckenna makes only the modest assertion that it is not the case that the agent is clearly not morally responsible in his case 6. If

¹ In what follows, I base my comments on Mckenna (2008).

the latter intuitive reaction attends case 6, then, given that Mckenna, similarly to Pereboom, has structured his cases in such a way that no morally relevant difference seems to apply to them, nothing prevents the intuitive reaction from generalizing to the remaining cases up to Pereboom's case 1 (Mckenna 153-54). The strategy Mckenna uses in reaching the intuitive result that it is not evident that Plum is not morally responsible in case 1 is to draw attention to the rich contingent of compatibilist agential factors in case 6. By fixing our attention on these factors, Mckenna then imports them in the remaining cases up to Pereboom's case 1. (152).

Mckenna concludes *that it is not the case that prof. Plum in case 1 is clearly not morally responsible* – a seemingly counterintuitive result, avoided by most compatibilists and incompatibilists alike, which gives Mckenna's argument its hard-line title. Moreover, Mckenna contends, his use of the generalization strategy against the four-case argument needs to elicit only the *indecisive* intuition that it is not clearly the case that Plum is not morally responsible in case 1. This would suffice to disable Pereboom's challenge to the compatibilist, for then it cannot be concluded, by way of generalizing from his case 1, that the normal case 4 is one in which it is *clearly* the case that Plum is *not* morally responsible (Mckenna 155).

It is in this sense that Mckenna's hard-line reply is supposed to establish the dialectical impasse between the two debating sides (Mckenna 154). The two conflicting intuitions about Plum's responsibility in case 1 are considered equally legitimate grounds for judgment. The lesson that Mckenna draws from the presence of this intuitive deadlock is that it amounts to a victory for the compatibilist since “she was only out to defeat an argument for incompatibilism and not to prove her compatibilist thesis” (154). Henceforth, I will assume that people can have the intuition that Mckenna seeks to elicit and investigate whether it can indeed be a part of the

intuitive evidence in the debate over the four-case argument.²

The indecisive intuition that it is not clear that Plum is not morally responsible in case 1, however, cannot be automatically taken as intuitive evidence, which one is to weigh against the opposing intuition elicited by Pereboom. The hard-line compatibilist should presumably take the indecisive intuition in question to misrepresent the fact of the matter in case 1 since by his own lights Plum is clearly morally responsible. According to the hard-line compatibilist, a commanding intuition that Plum is clearly morally responsible would be most appropriate. In the absence of some explanation of why the intuition that Mckenna can elicit is only an indecisive one, it is open to the hard incompatibilist to insist that this opposing intuition is a single-case aberration, an inexplicable oddity. Why should it be trusted with any evidential force if the hard-line compatibilist herself is bound to note that this intuition does not get the fact about Plum's responsibility straight?

In his own discussion on the subject in a 2008 paper, Mckenna offers a seemingly natural explanation of why proper concentration of our attention on all the satisfied agential conditions for moral responsibility is still not powerful enough to provoke the opposite intuition that Plum just *is* morally responsible. He offers the following considerations. The fact that our responses to manipulation cases are tested in an extremely different context – one that significantly differs from the contexts of ordinary practice out of which intuitions have evolved – explains why they are not commanding enough in the cases proposed by Pereboom (Mckenna 157).

Mckenna supports this assertion by highlighting some of our favorable reactions to more ordinary cases of manipulation which we encounter in daily life. He asks us to consider the case of a young child, Ann, who witnesses the slow deterioration of her mother's health, and

² I have not presented the specifics of Mckenna's six cases since I just assume in the discussion to follow that Mckenna can successfully elicit the intuition about case 1 which he needs. Limitation of space prevents me from doing justice to a full presentation of his cases.

ultimately, her death from leukemia. We are to suppose that Ann is of such tender age that her reasoning is not mature. Her early experience leaves an irremediable mark on her subsequent development so that she comes to see life as precious but limited – so vulnerable that it needs to be lived to its fullest. Later on, as a mature adult, Ann regards her experience as a real-life manipulation case – it is just that it is not manipulation by foreign agents but by the vagaries of life. Here, Ann's deeply entrenched values in life trace back to that decisive moment in her past when her mother was extremely sick. Moreover, she does not see that as an impediment to her freedom and responsibility but as a condition for it. Mckenna maintains that we should also regard it as such (156).

In essence, the case of Ann, construed as ordinary, real-life manipulation, is seen as a testament to our general ability to intuitively react to cases of manipulation in a favorable way. But if ordinary contexts of evaluation allow us to issue such judgments about cases of manipulation, then we can justifiably maintain that the salient difference between evaluative contexts (bizarre, thought-experimental ones vs. natural ones) is what ultimately explains the intuitive resistance to regarding Plum as clearly morally responsible in case 1. This seems to be Mckenna's argument.

II. Mckenna's Proposed Explanation: An Appraisal

Let us concede that there could, in principle, be some such real-life manipulation cases in which our responsibility-ascriptions are decisive. In other words, let us concede that we can still judge agents morally responsible for their actions even when they are manipulated in some more ordinary ways, by the vagaries of life, say. Does this suffice to establish the plausibility of

Mckenna's explanatory hypothesis? I believe there are strong considerations that go against it. Notice that his explanation of why our intuitions about moral responsibility are not commanding enough in extra-ordinary manipulation cases like Pereboom's assumes the following general form. Context A has provided all background conditions for the evolution and development of intuitions of type X. Context B significantly differs from context A. So if intuitions of type X are not as commanding in context B as in context A, given that all responsibility-relevant conditions are present in both contexts, it must be the case that the difference between A and B explains the indecisiveness of the intuitions. Because such an explanation assumes this general form, one would expect that it is applicable to structurally similar cases in which intuitions in the realm of moral responsibility are probed in vastly different contexts. This is not so, however. Mckenna's explanation falters when applied to the Frankfurt-style examples.³

Frankfurt-style examples famously involve much the same non-ordinary contexts for the testing of our reactions as do Pereboom's cases. However, no parallel skepticism about how commanding the intuitions elicited from Frankfurt-style examples can be, given the non-natural features of the situations, seems to apply here. In other words, no indecisive intuitions seem to be reported in these Frankfurt-style examples. It is not common, for example, to have the intuition that it is not clearly the case that Black is not morally responsible for killing Jones. Rather, these are usually quite decisive (which explains why PAP has more or less faded from prominence recently). It appears that the nature of the context – thought-experimental vs. a more natural one – does not preclude the generation of strong, *decisive* reactions in the latter cases.

One immediate objection calls for attention. It may not be clear that there is the necessary structural similarity between the two types of cases, which one can use to discredit

3 Harry Frankfurt introduces his famous examples in "Alternate Possibilities and Moral Responsibility," The Journal of Philosophy 66 (1969): 829-39.

Mckenna's explanatory hypothesis. For if we notice the particular details of these examples, a relevant difference in the function of the featured causal factors may well render the cases insufficiently alike. In a typical Frankfurt-style example it is the agent whose moral responsibility is probed that is assumed to be the most immediate causal origin of the action. The causal manipulation factor in Frankfurt-style examples is never embedded in the actual sequence of events but remains in the alternate sequence of events. The agent whose moral responsibility is tested in the imaginary cases is not manipulated in the actual sequence of events. It is not so with Pereboom's examples.

Here, one can attempt to explain away the decisiveness of the intuition that agents are morally responsible even if unable to do otherwise, which is the response to typical Frankfurt-style examples. Our intuitions, as the envisioned response has it, are commanding enough in Frankfurt-style examples, despite the unnatural contexts in which these are elicited, because the most unnatural features of these contexts, the causal manipulation factors, are not as directly relevant to our intuitions about the case. These unnatural features are, in a sense, only background characteristics of the case. This is because the causal manipulation factors are not embedded in the actual sequence of events that leads to the agent's action. The unnatural evaluative contexts of Frankfurt-style examples do not interfere with our normal intuitive responses because these contexts' most eccentric features, the causal manipulation ones, do not so interfere with our intuitions.

Note, however, what the foregoing explanation implies for Pereboom's cases. Mckenna's original proposal for an explanation of why his newly elicited intuition about case 1 can only be an indecisive one has to be somewhat modified. It will not only be the case that the mere unnaturalness of the evaluative context explains the indecisiveness of the intuition that it is not

clear that Plum is not morally responsible. According to the modified proposal, the features of the context that are most unnatural directly interfere with our normal intuitive responses. This would be what ultimately explains why our intuitions are wavering with respect to case 1. But let us ask, what is the prominent, most eccentric or unnatural feature of Pereboom's case 1? Presumably, this is the causal factor, Plum's manipulation by neuroscientists. This is the factor that we characterized as the background, non-obstructive one in the Frankfurt-style examples - that is, non-obstructive for our normal intuitive responses. We should stress, then, that on the present proposal the indecisiveness of the intuition that Mckenna reports can be best explained by the presence of this unnatural causal factor, which interferes with one's normal intuitive responses and, arguably, obstructs one's proper attention to the relevant agential factors.

What is problematic about the Mckenna's modified explanatory hypothesis is that it cannot provide any account of why the exact same configuration of agential and causal factors in case 1 – the neuroscientists' manipulation of Plum and the rich compatibilist agency conditions for moral responsibility – generates a very different response in incompatibilists and soft compatibilists alike.⁴ Their reaction is, in effect, that the example is one of clear *non-responsibility*. The presence of unnatural contextual factors of the causal variety, Plum's manipulation by neuroscientists, construed as those unnatural features of the context that interfere with normal intuitive responses, is supposed to explain the indecisiveness of Mckenna's considered intuition. But why is it that precisely these same obstructionist causal factors do not so interfere with other compatibilists' clearly opposing intuitions that Plum just *is not* responsible

4 Soft-line compatibilism denotes the position that gives the so called soft-line reply to the four-case argument. This kind of reply is, for example, given by Al Mele who in fact seeks to affirm the intuition reported by Pereboom (that Plum is clearly not morally responsible when manipulated in case 1) but give it an alternative explanation. One common compatibilist response to the case is that the presence of manipulation, rather than causal determination, is what explains the non-responsibility intuition. See Mele's "Free Will and Luck: Reply to Critics," *Philosophical Explorations* Vol. 10, No 2, (2007) p. 202-206.

for murdering Ms. White in case 1?

In sum, we have a bad dilemma. If we accept Mckenna's original explanation, according to which the difference in unnatural vs. ordinary contexts of evaluation explains the indecisiveness of the intuition, we will have to arbitrarily restrict its scope so that it does not apply to Frankfurt-style examples, despite the similarity between the cases. If, on the other hand, we accept what seems to be the only plausible way to modify Mckenna's original proposal in response to that difficulty, we will not be able to account for why most compatibilists' intuitions are decisive and not obstructed by the presence of any unnatural causal manipulation factors in the case.

However, there is one way in which one could escape the foregoing dilemma. It is open to the compatibilist to consistently reject any theoretical use of unusual evaluative contexts, including the ones featured in Frankfurt-style examples. Such a compatibilist position might claim support from the ordinary practice of responsibility-ascription as it has evolved in our communal life. According to this position then, we should limit our attention to the patterns of evaluative judgment that are elicited from ordinary language use.⁵ The lesson to draw from Mckenna's indecisive intuition is precisely that unusual evaluative contexts are unreliable guides to people's fundamental views on moral responsibility.

There are at least two main difficulties with such a position. First, it might have to assume that the apparent lack of uniformity and coherence in people's ordinary responsibility-ascriptions can be addressed without any resort to more unusual examples. But there is good reason to doubt that this could be accomplished. It is widely acknowledged that when people are faced with causal-historical considerations of *no* extra-ordinary variety they often retract their initial intuitive charges. Early-childhood manipulation and harsh living conditions which might

⁵ This reply was suggested to me by Gordon Barnes.

incline individuals toward criminal behavior usually make people reconsider their initial intuitive reactions and make them more lenient toward such individuals. But even when the causal-historical factors are not of such a prominent sort, people still seem less inclined to judge others when they know more about the evolution of their character and behavior rather than about some small time-slice of others' lives. It is of course debatable whether all this favors an incompatibilist interpretation of these patterns of ordinary thought. It could be the case that the causal-historical factors in question just affect people's judgment as to whether the proper *agential* conditions of responsibility such as reason-responsiveness, rational self-control, etc. actually obtain in these cases. A bad personal history might make a bad compatibilist agential structure.

The point here is that the role of unusual evaluative contexts like the ones featured in the four-case argument is precisely to make both agential and causal-historical factors equally prominent and hence test which ones prevail in people's judgments. One could thus have some hope of sorting through the multitude of considerations that might determine how people respond to cases in *ordinary* contexts. It is hard to see how one can consistently reject the use of atypical evaluative contexts and be able to discern solely through people's patterns of responses in ordinary language whether their intuitions are ultimately compatibilist or incompatibilist in nature. This is because ordinary ascriptions of responsibility seem to be of a mixed character. What is needed is "controlled" settings, which in turn require unusual circumstances.

There is another difficulty with the rejection of atypical evaluative contexts, however. The incompatibilist can rightly insist that the evolution of our practice of responsibility-ascriptions could be biased against any reflective consideration of causal factors. After all, the demands of a successful communal life might have been such that the members of a given

community were more concerned with controlling the moral character of the other members than anything else. This could have been served best if the practice of responsibility-ascriptions emphasized the presence of agential traits such as rational self-control, receptivity to moral reasons, etc. Hence a restriction to merely ordinary contexts is bound to favor agential compatibilist considerations. This should be seen as objectionable since the incompatibilist can rightfully maintain that when people's judgments are not interfered with by the pressure of custom, people can be properly sensitive to causal-deterministic factors. In fact, the idea of the four-case argument is precisely to make these factors salient so as to bring about that sensitivity (Pereboom 100). Restricting the relevant test cases to the ordinary ones does not give a fair hearing to the incompatibilist position.

I take it that no convincing explanation of why we can only have indecisive intuitions about manipulation cases like Pereboom's case 1 seems forthcoming. In the absence of such an explanation, the evidential force of Mckenna's newly elicited intuition, in particular, seems substantially undermined. One cannot automatically assume that Mckenna's indecisive intuition can counterbalance the commanding intuition of incompatibilists and soft-line compatibilists just in virtue of its availability. The hard-line compatibilist has to acknowledge that the indecisive intuition that it is not clear that Plum is not morally responsible misrepresents the truth of the matter in case 1 since by the hard-line compatibilist's lights Plum is clearly morally responsible. In the absence of any good explanation of why one cannot have a commanding intuitive response to that effect, given that such a response would be most appropriate, the intuition cannot do the evidential work that the hard-line compatibilist needs it to do.

III. Dialectical Burdens

In the preceding section, I defended two central claims. I argued that the hard-line compatibilist cannot assume that the newly elicited indecisive intuition about Pereboom's case 1 does any evidential work without a proper explanation of why it is merely an indecisive one. I further argued that Mckenna's proposed explanation is inadequate as it stands. In the absence of any positive support for her reliance on this intuition in countering the four-case argument, the hard-line compatibilist cannot successfully answer the challenge posed by the argument.

However, there is one last difficulty in accepting the foregoing conclusion that we need to attend to. This difficulty pertains to matters of dialectical burdens. It might be objected, in light of some of Mckenna's other remarks, that I have not made a sufficiently good case for the demand that the indecisive intuition be given an explanation. My requirement for such an explanation can be seen as ill-founded because the hard-line compatibilist in the face of Mckenna has discharged his dialectical duty, so to speak, simply in virtue of eliciting a contrary intuition about case 1. All that is needed in order to counter Pereboom's argument is that one demonstrate the weaker claim that it is not clearly the case that Plum is not morally responsible in case 1. The incompatibilist is the one who advances the claim that "appropriately manipulated agents" are not morally responsible. She is the one who started the confrontation, as it were. Thus, according to Mckenna, the compatibilist's dialectical burden is a comparatively lighter one. He only needs to demonstrate that this is not clearly the case, which he does, in Pereboom's case 1, by eliciting an indecisive intuition to that effect. As Mckenna puts it, "If the compatibilist can merely cast doubt on the incompatibilist's positive claim, then she has adequately replied to the argument *even if she has not offered a positive defense of her thesis regarding manipulation*" [italics in

original] (Mckenna 155).

Mckenna's thesis is that an agent who is manipulated in such a way that CAS is satisfied is free and morally responsible. As he makes it clear, he does not intend to prove this thesis. In line with his remarks on the matter of dialectical burdens, he has assumed that he only needs to demonstrate that a properly manipulated agent like Plum is not clearly absolved from moral responsibility. This is achieved by garnering contrary intuitive evidence. But if this is the case, then our demand that he offer some explanatory account of his newly elicited intuition about case 1, on pain of leaving it evidentially suspect, will seem illegitimate. It presupposes that the burden of proof is on the hard-line compatibilist's side. True, to provide an explanation of why people's intuitions cannot commandingly attest to the responsibility of properly manipulated agents, even though this would be the most appropriate judgment, according to hard-line compatibilism, is not to provide a positive *argument for* the latter position. It is, nevertheless, tantamount to offering some defense of the thesis that manipulated agents who satisfy CAS just *are* free and morally responsible. For a possible explanatory account of our wavering intuitions in cases like Pereboom's case 1 will help to deflect a readily available criticism of Mckenna's thesis, namely that people's intuitive reactions do not seem to conform to that thesis. Accordingly, our demand for an explanation of the indecisiveness of Mckenna's newly elicited intuition amounts to a demand that Mckenna offer some positive defense of his thesis, in some broad sense of defense. Yet, Mckenna contends that the hard-line compatibilist has discharged his dialectical duty. We seem to be facing a difficulty. What should we say about the matter of dialectical burdens in this case? Is it sufficient for Mckenna to simply elicit an opposing intuition about case 1 in order to counter the four-case argument, contrary to what I have been arguing in the preceding chapter?

The best way to approach these questions is to examine Mckenna's grounds for

maintaining that his six-case argument has shifted the burden of proof to the incompatibilist's side. As already noted, Mckenna claims that the compatibilist's dialectical burden is “comparatively lighter” than the incompatibilist's. What Mckenna's somewhat cursory remarks suggest is that he takes the fact that the incompatibilist is the one advancing a manipulation argument to be a sufficient reason to counter his argument without invoking any considerations in support of his own hard-line position. Merely eliciting the opposing intuition that Plum is not clearly absolved from moral responsibility should thus suffice for shifting the burden of proof back to the other side.

This should not satisfy us. If the above exhausts Mckenna's reasons for claiming a discharged dialectical duty, then it is markedly inadequate. Arguably, judgments about dialectical burden should be sensitive not just to balance sheet considerations, who started which argument, as Mckenna seems to assume. Rather, the nature of the arguments and claims being advanced should surely be weighed in the determination of argumentative burden. In particular, it seems that one ought to consider how much of commonsensical thought hangs on the truth or falsity of the claims in question. How revisionary are they with respect to widely accepted belief in the domain under consideration? What other beliefs do they bear on and in what way? Such considerations about a given claim's content, whether pragmatic or theoretical, are surely relevant to how much argumentative support ought to be mustered for this claim.

One plausible line of thought from some of the epistemological literature is that the justificatory status of a belief should reflect how much care a believer has taken to attend to the evidence for and against her belief. Arguably, how much or how little evidence ought to be attended to will depend on how much hangs on that belief, i.e. on the consequences of adopting the belief in question. Simon Blackburn has articulated the thought particularly well: “An

interesting consequence of this account [Blackburn's own position on reasonable belief] is that whether it is reasonable to have confidence in a proposition may sometimes depend upon what we are going to do with that confidence...For if the consequences of our having misplaced confidence in a proposition are anticipated to be rather serious, then it may be unreasonable for us to fail to collect evidence, which in a less serious situation, it might have been reasonable to fail to collect” (Blackburn 36).

In other words, our judgment about how much evidential support a given belief warrants seems to be rightly sensitive to the consequences of holding that belief. In the case of theoretical belief, these consequences will be determined by, among other things, the content of that belief, that is how much it forces changes elsewhere in the relevant doxastic system. These consequences will thus determine how revisionary (if at all) this belief happens to be with respect to the larger doxastic system in which it is featured. This is a plausible thought. It seems that the wider the influence of a given belief on others, especially in forcing changes, the greater the evidential support that is required to count as reasonably holding that belief.

If the reasonableness of beliefs partially depends on how much evidential support one has for these beliefs, and if the necessary support is sensitive to considerations about how much hangs on these beliefs, then it seems we have some reason to extend these considerations to our present case. The success of the hard-line compatibilist’s effort to discharge his dialectical duty will not depend merely on his ability to gather opposing intuitive evidence to bear on Pereboom's case 1. Rather, the nature of the conclusion of his six-case argument will have to be weighed in the calculation of argumentative burden. If this conclusion has a widely counter-intuitive, revisionary character, then Mckenna cannot assume his task to be accomplished simply in producing this negative argument.

How much positive evidential support (if any) the conclusion of Mckenna's his six-case argument needs will depend on its conservativeness or lack thereof. If such a contention seems contestable, then we might relativize the conservativeness in question to the framework of ordinary compatibilism. If, in order to counter Pereboom's four-case argument, Mckenna finds himself endorsing a conclusion that seems too revisionary by most compatibilist's own lights, then this conclusion had better receive some positive support, in the broad sense that I have already noted. Furthermore, and this is the key claim here, if the above is correct, then Mckenna cannot assume to have shifted the burden of proof to the incompatibilist side simply in virtue of producing some counter-intuitive evidence that supports an especially revisionary conclusion.

That the conclusion of Mckenna's six-case argument is revisionary seems hardly disputable. Recall that Mckenna claims to have countered Pereboom's four-case argument by using the generalization method to establish the conclusion that it is not clear that Plum is not morally responsible in case 1. Given that Plum is the quintessentially manipulated agent who the majority of compatibilists do not deem morally responsible, Mckenna's conclusion breaches a central commitment of compatibilism. This is the familiar commitment that the absence of manipulation in the causal production of a given action is a necessary condition for being responsible for that action. To assert even the weaker claim that it is not clear that an agent like Plum is not morally responsible is already to be making a revisionary claim in light of compatibilism's own theoretical framework. The consequences of accepting Mckenna's conclusion for compatibilism in general are quite serious for they force a significant modification of traditional compatibilism.

Furthermore, Mckenna seems to be sensitive to the revisionary nature of his conclusion for he immediately goes on to soften its effect by trying to tell a story about ordinary judgment in

cases of familiar manipulation. This is where his presentation of the story about Ann enters into his argument. Mckenna thus gives the impression that he is concerned to place his counter-intuitive conclusion in a better context, that is, to stress that it may not be as radical a departure from the ordinary practice of responsibility-ascription as it may first appear. He seems to be well aware of the theoretical pressures toward conformity to a standard of conservativeness.

It is important to notice what is not being argued here. The foregoing considerations do *not* target Mckenna's claim that the hard-line compatibilist need not prove his general thesis about agents manipulated in an appropriate way in order to effectively respond to the four-case argument. This could be granted for the sake of argument. Yet, the weaker conclusion of his six-case argument does need some positive defense in light of its revisionary nature, if the compatibilist is to discharge his dialectical duty. As already noted, the demand for such a defense here takes the form of an expectation for an explanatory account. Mckenna needs to provide some explanation of the indecisiveness of the opposing intuition about Pereboom's case 1 in order to forestall an immediate objection one may have to it, namely that it is an aberration that we ought not to take into consideration. It is easy to see that this objection is motivated by the counter-intuitive and thus revisionary nature of Mckenna's conclusion, which he bases on this indecisive intuitive evidence. In this respect, the main thrust of our discussion so far has been that matters of dialectical burden are not to be decided without proper sensitivity to the nature of Mckenna's conclusion.

It seems that the only way that one can resist the force if these considerations will be to argue that even if a positive defense of the counter-intuitive conclusion of the six-case argument is needed, it is not needed for the purposes of answering the incompatibilist challenge in the debate over the four-case argument. One might maintain that offering such a positive defense is

optional. It may be required to address the concerns of fellow compatibilists but it remains a separate issue that does not need to be resolved in order to counter the four-case argument, nonetheless. There is a danger here that such a line of thought will license a questionable general policy for answering the criticism of one's opponents. Whenever an argument is being advanced against one's position, one might counter it by endorsing some deeply counter-intuitive and revisionary conclusion, championing the cause of hard-line responses, and relegating the positive defense of it to other occasions. In such a case, the defense of one's claims will not be owed to the opponent but to one's in-group.

In the debate between compatibilism and incompatibilism in which all sides are typically concerned with mustering intuitive support for their claims, getting away with hard-line responses may not be a good policy to recommend. It seems that these ought to be held to a high standard of warrant. This is especially relevant to our present concern with the matter of dialectical burdens. Shifting the burden of proof to the other side so easily, without allowing the nature of one's conclusion to be weighed into the calculation and distribution of argumentative responsibility, is surely a sign of leniency. We should not be satisfied with the hard-line compatibilist's grounds for rejecting demands for some positive support of his conclusion. However, if these remarks do not seem convincing, then we could even grant that for the purposes of answering the four-case argument Mckenna does not need to provide a positive defense of his conclusion. Yet, given its revisionary nature, we are free to insist that it immediately generates a burden of proof on the part of the hard-line compatibilist so that we are still owed some defense of his indecisive conclusion. The dialectical burden is not on the side of the hard incompatibilist, however we decide to frame the issue. I take it that this suffices to make the case that our demand for some explanatory support of Mckenna's conclusion is not

unjustified. In light of the arguments from the previous section, this demand has not been met.

IV. A Diagnosis

Yet, one could still harbor some concerns about the general state of the debate that emerges from the success of Mckenna's generalization strategy. Even if the indecisive intuition that Mckenna elicits does not carry evidential force, there is still the fact that Mckenna can elicit this opposing intuition. What are we to make of this fact? We are assuming that Mckenna can successfully elicit a conflicting intuition about Pereboom's case 1 using the very same method of generalization that Pereboom uses. Mckenna's starting point, his normal sixth case, featuring an agent who is determined by the past and the laws of nature, is different than Pereboom's starting point, case 1. What might be seen as puzzling here is that one's different starting point, whether it will be a clear-cut manipulation case like Pereboom's case 1 or Mckenna's ordinary case 6, together with a pressure toward generalization, can yield different intuitive results. Perhaps, what really needs to be explained here is the role of the starting point from which one generalizes. Moreover, this would appear to be more of a general problem for both theoretical positions than merely a problem for the compatibilist or incompatibilist side.

The compatibilist would presumably maintain that a normal scenario like Pereboom's fourth one, in which natural deterministic processes lead to Plum's murder of Ms. White, is bound to yield a different intuitive response in ordinary circumstances, that is, outside the context of the other four cases that Pereboom presents. The compatibilist is likely to insist that in ordinary circumstances people will unwaveringly judge Plum to be morally responsible in case 4. Arguably, the compatibilist should also allow that at least some ordinary people might

have conflicting intuitions about case 4 when they are initially subject to Pereboom's generalization method. To suppose otherwise might be to have an unreasonable level of confidence in ordinary people's judgments by taking them to be unwavering through and through. So I take it that it should be puzzling even for the compatibilist how a different starting point together with generalization pressures can alter people's intuitive responses to a single case.

Notice that, in light of our previous discussion, one cannot simply resort to the distorting influences of unnatural contexts in explaining away people's conflicting intuitions. In other words, one cannot argue that it is not really the influence of the variable *directions* of generalization pressures from distinct starting points that yields the different intuitive results in the examples at hand but these other distorting influences. Furthermore, it is not open to the compatibilist and incompatibilist alike to give an alternative explanation that cites the disparate influence of causal-historical and agential factors on people's reactions to the starting-point and end-point cases. It has already been emphasized that both kinds of factors are taken to be equally prominent in all cases. There seems to be at least a *prima facie* case for demanding a distinct explanation of the role that the starting points in series of cases play in people's different intuitive responses.

In order to give a possible account of the mechanism that gives rise to conflicting intuitions in the debate over the four-case argument it may be instructive to first consider the role of consistency pressures on responsibility-ascriptions. One's particular evaluation of the first case in a given series already generates some pressure to give a similar evaluation to the second one if no responsibility-relevant difference seems evident. This pressure is then carried over to the remaining cases. How might such an observation help with explaining why single cases like Pereboom's first and last one can yield different intuitive results depending on whether the case

considered first is a normal one or a clear-cut manipulation case?

One possibility here is that consistency pressures simply have different directions depending on *the starting point* in question. What gives rise to the consistency pressures is the theorist's deliberate construction of a series of cases in which the members do not seem to differ in any responsibility-relevant way. If one is presented with a normal case featuring naturalistic determination, for example, on the condition that the successive cases in the series seem to feature the same compatibilist-friendly agential structure, then the pressure to be consistent with one's initial intuitive judgment in the normal case might lead one to waver on the clear-cut manipulation case. Conversely, if one is first presented with a clear-cut manipulation case and subsequently with a series of cases that do not seem to differ with respect to the causal-deterministic factors featured in them, then one might feel the pressure to be consistent with one's initial intuitive reaction to the clear-cut manipulation case. Given that most people will have an especially firm judgment that the manipulated agent is not morally responsible and given that the direction of consistency pressures is different in this case, one might well end up more inclined to extend the same judgment to a normal case of naturalistic determination. This mechanism may then allow one to make some progress in explaining the different intuitive results of Mckenna's and Pereboom's use of the generalization strategy.

What emerges from the foregoing discussion is that the nature of the case that is considered first, whether it will be an ordinary one featuring regular deterministic processes or a clear-cut manipulation case, makes all the difference. This is the sense in which the different *direction* of consistency pressures, which operate from a distinct starting point in one's evaluation of a series of cases, determines the conflicting intuitions that one ends up with. Perhaps the lesson to draw here is that not all starting points are equally reliable. The debate

between the compatibilist and incompatibilist will then involve arguments to the effect that the opposing side's starting point is not a legitimate case from which to generalize.

I have argued that the compatibilist cannot disqualify the incompatibilist's atypical manipulation cases as illegitimate on the grounds that they feature contexts that are too extraordinary and so lead people's normal intuitions astray. Nevertheless, it is open to the compatibilist to argue on some other grounds that clear-cut manipulation cases like Pereboom's case 1 are not reliable starting points in generalization arguments. In a similar vein, it is open to the incompatibilist to argue that something about cases like Mckenna's sixth one and Pereboom's fourth one, which feature only naturally determined agents, makes them unreliable starting points. The conflicting intuitions that the different starting points together with generalization pressures yield thus seem to reflect a general problem that arises from taking all cases to be equally reliable. Perhaps, such an assumption will need to give way.

The foregoing discussion is not intended as any definite account of what generates the puzzling state of the debate between the hard-line compatibilist and the incompatibilist. Yet, if progress is to be made beyond the purported dialectical stalemate between the two debating sides, then it seems that one ought to consider the full scope of possible explanations of the particular mechanism that operates when people make responsibility-ascriptions. People's intuitions seem particularly sensitive to whether a given case is evaluated first in a series of cases, in isolation from others, or last, in the end of a process of generalization. Proper attention to such observations may enable one to advance the debate by locating the disagreement between the two sides in the particular application of generalization strategies. What kind of cases are to serve as reliable starting points from which to generalize? What are the criteria on which one is to judge what makes a case good for the purpose? These questions are worth exploring as they

may allow one to make some progress in the debate over the four-case argument.

V. Advancing the Debate

In chapter four, I argued that proper sensitivity to the conflicting results yielded by the use of the generalization method in the debate over the four-case argument should lead us to a certain conclusion. I suggested that the lesson to draw from the current state of the debate is that not all kinds of responsibility-relevant cases should be treated as equally reliable starting points in our evaluation of a series of examples. In light of these developments, I maintain that the disagreement between the compatibilist and incompatibilist should be located in the particular application of the generalization strategy. In order to advance the debate we should thus focus on Mckenna and Pereboom's starting points – ordinary causal-deterministic cases versus causal-manipulation cases. Does the hard incompatibilist have any reason to criticize the use of Mckenna's ordinary causal-deterministic example as a starting point for evaluating the series of cases leading to Pereboom's case 1? Does she have any good argument to the effect that the generalization strategy needs to be confined to the kind of use she puts it to, namely a generalization of intuitive results *from manipulation cases*?

The central contention of this chapter is that the hard incompatibilist does have serious grounds on which to object to Mckenna's use of ordinary causal-deterministic examples as starting points from which to generalize. The discussion to follow is one extended argument for this claim. If I am correct in what argue, then the threat posed by Mckenna's hard-line response to the four-case argument stands effectively annulled as this response depends on the legitimacy of the starting point from which he generalizes.

There are at least two kinds of objections that one might have to the hard-line compatibilist's use of the generalization strategy – methodological and more substantive ones. I will first address the former kind. One immediate objection that one can levy against the six-case argument is that it takes as its starting point the sort of case that is bound to evoke different intuitive reactions on the part of the compatibilists, on the one hand, and incompatibilists, on the other. But if the generalization method has proved to yield conflicting results about Pereboom's case 1, depending on whether this case is evaluated first in isolation from the other ones or last after Mckenna's other cases, then we seem to have reason for putting greater confidence in the less divisive case, namely Pereboom's case 1.

We have abandoned the assumption that all starting points are equally reliable in light of such conflicting results. The natural reaction seems to be to put greater confidence in the starting point case that manages to generate more agreement in the intuitive judgments of those who evaluate it. Cases featuring regular causal-deterministic processes like Mckenna's sixth one are hardly ones that generate such agreement. In contrast, the strength of Pereboom's argumentative strategy seems to lie precisely in the choice of a starting point case that generates agreement between incompatibilists and most compatibilists alike. The presence of causal manipulation factors is intended to make the compatibilist agree with the judgment that Plum is not morally responsible in case 1. The incompatibilist presents an example from which to generalize that is by and large not divisive. It is thus open to the hard incompatibilist to argue that Mckenna's sixth case is comparatively more divisive than Pereboom's case and thus ought not to serve as a starting point from which to generalize. Given that both sides are now forced to have a more discriminatory approach to the use of the generalization strategy, in light of our preceding arguments, the most straightforward criterion for discrimination is surely how much antecedent

and broad intuitive force a given case can marshal.

As plausible as this initial argument may be, it is unlikely to be sufficiently persuasive to the hard-line compatibilist. Mckenna may after all insist that he has a perfectly good justification for relying on a comparatively more divisive case to start off the generalization process. This is the already familiar thought that he is only responding to an argument and hence not out to convince his opponent to adopt his own position. The requirement that one start with the kind of responsibility-relevant example that can marshal the intuitive support even of one's opponents is appropriate only in the event that one is advancing a positive argument that has as its purpose to convince others. The function of Mckenna's own six-case argument is not to convince the incompatibilist that it is not clear that properly manipulated agents are not morally responsible. The imposition of a requirement that this sixth case have a level of intuitive support comparable to that of Pereboom's case 1 is entirely illegitimate. But then the fact that the sixth case generates less intuitive agreement is not to be invoked as a ground for discriminating against it. If Mckenna's application of the generalization strategy is to be cogently criticized, then the feature that renders Mckenna's sixth case a less reliable and objectionable starting point must be sought elsewhere.

In light of these considerations, which may or may not be conclusive, the hard incompatibilist should instead seek a more secure ground on which to disqualify Mckenna's sixth case as a reliable input for the generalization mechanism. The hard incompatibilist needs a more substantive argument for this purpose. The difficulty in offering such an argument is especially compounded by the fact that Mckenna's sixth case is so structured as to feature *both* agential and causal factors. The natural deterministic causes operating in Plum's decision to kill Ms. White are in fact brought to the surface. This makes the following tempting line of thought powerless to

support the rejection of normal cases as starting points. Initially, one may be tempted to say that natural, manipulation-free cases cannot reliably serve as inputs to the generalization process because they are representative of the vast majority of cases which people encounter in everyday life. Crucially, these are cases in which people are not generally aware of the causes of their actions. Any application of the generalization method to such scenarios is thus bound to be objectionable to the hard-incompatibilist. Yet, as already noted, Mckenna's sixth case evades this possible criticism since it features the deterministic causes operative in the agent's decision to kill.

The possibilities for a successful argument on the part of the hard incompatibilist seem to be significantly circumscribed. Moreover, since the hard incompatibilist has traditionally relied on the Spinozistic intuition that our illusion of freedom is maintained by an ignorance of causes and that this illusion would be dispelled if we were aware of these cause (Pereboom 2005; Mckenna 2008), a special difficulty emerges for her argument. Not only will it seem that normal, manipulation-free cases can be reliable starting points but it will also seem that the Spinozistic intuition can thereby be substantially weakened. If, as we have been assuming for the purposes of this paper, people can share Mckenna's indecisive intuition about his sixth case, despite the clear presence of the deterministic *causes* behind Plum's action, then the Spinozistic intuition seems to be ineffective in generating incompatibilist sympathies. More importantly, the traditional reliance on this Spinozistic intuition happens to operate *against* the incompatibilist in the present case as it creates the expectation that people would in fact have the appropriate incompatibilist intuition about Mckenna sixth case so that there could be *no* good objection to using this case as input to the generalization process. Pointedly, it is the incompatibilist's confidence in our ability to be properly sensitive to the revelations of the ultimate Spinozistic

perspective that creates the most difficulty for advancing an argument against Mckenna's application of the generalization strategy.

What grounds for rejecting Mckenna's manipulation-free case are then left to the incompatibilist?

This can become clear only after she draws the correct lesson from the failure of the Spinozistic intuition to support her present case. The hard incompatibilist ought not to have placed such high confidence in the Spinozistic suggestion in the first place. The thought that our illusion of freedom would be dispelled if we became aware of the causes of our actions because it is ignorance of these causes that maintains it should preserve an important place in incompatibilist theorizing. I do not intend to challenge that. Yet, the mistake that the incompatibilist has made, manifested in her present failure to resist Mckenna's application of the generalization method, is not to recognize the limitations of that thought. It should be acknowledged that mere awareness of the causes of our actions may not suffice to incline the majority of us to react in an incompatibilistically-friendly way. This is because ignorance of causes is unlikely to be the central basis for people's illusion of freedom, though it admittedly has a significant role to play in maintaining this illusion. Once it is recognized that other powerful factors generate this illusion, which are, moreover, operative in our responsibility-ascriptions in normal, manipulation-free cases like Mckenna's sixth one, it should become clear why the incompatibilist can reject these latter cases as unreliable starting points.

I suggest that the enduring force of the illusion of freedom does not stem primarily from ignorance of the causes of our actions but from the nature of the deliberative standpoint essential to the exercise of agency. Following insightful remarks by F.P. Ramsey on the critical role of the deliberative standpoint in causal reasoning (Ramsey 146), I want to suggest that the very nature

of our deliberative standpoint is the main source of the potentially wavering response people may have to ordinary manipulation-free cases, even when the ultimate causes of action in these cases are prominently displayed. If this is indeed the case, then the hard incompatibilist will be able to explain why it is still possible for some people not to have the clear incompatibilist intuition to Mckenna's sixth case, when the latter is evaluated first, despite the presence of deterministic causes in the case.

On this view, the illusion of freedom is inescapably embedded in the deliberative standpoint and it will be implausible to expect that it can be so easily dispelled when people become aware of the ultimate deterministic causes of their actions. By extension, it will be implausible to expect that an awareness of the natural deterministic causes of others' actions will also easily and universally incline people to unwavering incompatibilist sympathies. There are bound to be great pressures to sustain the illusion of freedom. The incompatibilist will thus have some good grounds to resist the use of ordinary cases as starting points in the application of a generalization strategy. It is only by introducing manipulation factors, which can highlight the real import of deterministic causation, that one can hope to truly test the intuitions of people. The generalization method should first be applied to cases that allow people to resist the overwhelming pressures to maintain the illusion of freedom. The rejection of ordinary cases as unreliable will be justified in terms of their inability to counterbalance the inherently strong pressure to imagine our actions as “an ultimate contingency.” Otherwise, the incompatibilist side will not have received its fair hearing.

Ramsey's key insight is that the basics of causal reasoning emerge from the deliberative standpoint: “Again from the situation when we are deliberating seems to me to arise the general difference of cause and effect” (Ramsey 146). It is from the perspective of our own agency that

we understand the concepts of cause and effect. But given such a fundamental relation between causal reasoning and the deliberative standpoint, one would expect that certain characteristics of this standpoint will have an essential impact on one's understanding of the causes of one's own action. This is precisely what Ramsey's remarks further reveal: "In a sense my present action is an ultimate and the only ultimate contingency" (ibid.). Now, one should resist the implied claim that no contingency in the action of others is recognized. For, barring skepticism about the inner life of others, it seems that whatever feeling of contingency in action one may feel from one's own deliberative perspective, one will thereby transfer in the evaluation of others' actions. But if that is the case, then it seems that it is integral to the normal exercise of agency and recognition of others as agents that one recognize the ultimate contingency in one's actions and the actions of others.

Furthermore, it becomes clear that any expectation that this belief in ultimate freedom should be universally and easily dispelled once the deterministic causes of others' action are revealed is ill-founded. One would expect, on the contrary, that such a fundamental belief, constitutive of the deliberative standpoint and thus the very nature of causal reasoning (if Ramsey is to be trusted on this score) should be hard to counteract. Given this difficulty, the hard incompatibilist already has the rudiments of a case against the reliance on ordinary cases as starting points from which to generalize. For these will not be sufficiently removed from the familiar background against which ordinary causal reasoning operates and thus biases one towards a presumption of ultimate contingency in action.

As already intimated, one's evaluation of a given agent, who is assumed to engage in relevantly similar deliberative processes as one engages in, is constrained by the presumption of "contingency." This presumption is the centerpiece of our ordinary conception of agency, which

informs our common interaction with others. It stands to reason that such a presumption ought to be counterbalanced for the purposes of probing the hard incompatibilist case against moral responsibility. The key claim then is this. The introduction of manipulation in the causal production of Plum's action can accomplish an essential task for the fair evaluation of Plum's responsibility in Pereboom's case 4. The essential task is to remove any biased tendency toward a compatibilist evaluation of normal agents that is, crucially, a product of the ordinary pressures to attribute contingency to the actions of others *just* in virtue of their being deliberating agents like us. This residual source of pressure, inherent in our conception of agency, needs to be counteracted. The incompatibilist can thus insist that if one chooses to generalize from normal cases, then one thereby allows these pressures to unfairly bias one's intuitions against the incompatibilist. The presence of manipulation, on the other hand, can make the presence of causal determinism stark enough so as to act against these pressures so the compatibilist/incompatibilist case can be judged in abstraction from these pressures. Therefore, Mckenna's ordinary sixth case ought to be rejected as a starting point of generalization.

Now the hard-line compatibilist might still insist that there is no good reason to dismiss the presence of natural deterministic causes in Mckenna's sixth case as inadequate for the task of removing whatever fundamental pressures to maintain an alleged illusion of freedom people might to be subject to. Wouldn't highlighting the deterministic causes of Plum's action in the sixth case undermine these pressures if people were indeed incompatibilists at heart? Contrary to what has been said so far, it might seem reasonable to suppose that the causal factors featured in Mckenna's case can do the job. But the incompatibilist is bound to disagree. One should not expect that mere awareness of deterministic causes can act against such a deeply ingrained presumption of contingency in action implicated in our very conception of agency. For it is quite

evident in the common practice of responsibility-ascription that even after individuals soften their reactive attitudes in response to increased knowledge of the causes behind others' actions, they do not completely absolve them of responsibility. It is common to encounter individuals who readily enumerate all the conditioning that criminals have undergone in their childhood, the harsh life conditions, abuse, neglect, etc. that these criminals have been subject to and still respond after all that “they were nonetheless responsible.”

Notice that this does not necessarily testify against incompatibilism as such reactions are hardly a universal phenomenon. However, what they do show, according to the incompatibilist, is that it is unlikely that mere ignorance of causes fosters the illusion of freedom. One cannot expect that even after becoming aware of these causes individuals will completely abandon the illusion of freedom for its source is a deeply ingrained presumption of contingency in action which is integral to our common conception of agency. Given such a fundamental source, the incompatibilist ought to insist that normal cases like Mckenna's are unlikely to fairly address the need for counteracting this presumption in order to give the incompatibilist side a hearing.

Conclusion

I have argued that the alleged dialectical stalemate between the hard incompatibilist and compatibilist sides in the debate over the four-case argument does not in fact hold. Michael Mckenna's hard-line reply does not succeed in countering the intuitive evidence garnered from the argument. I have further offered a diagnosis of the current state of the dispute between the two theoretical positions. Variable directions of consistency pressures in the divergent application of the generalization method have led to opposing intuitive results in the four-case and six-case arguments. I have argued that the debate can be advanced once the particular application of the generalization method becomes the locus of disagreement between the two

sides. On this basis, I have offered an argument in favor of the hard incompatibilist use of the generalization method, which is supposed to vindicate the four-case argument.

Works Cited

- Blackburn, Simon. *Reason and Prediction*. Cambridge, UK: Cambridge UP, 1973. Print.
- Fischer, John Martin, and Mark Ravizza. *Responsibility and Control: A Theory of Moral Responsibility*. Cambridge, UK: Cambridge UP, 1998. Print.
- Frankfurt, Harry. "Alternate Possibilities and Moral Responsibility." *Journal of Philosophy* 66.23 (1969): 829-39. Print.
- Frankfurt, Harry. "Freedom of the Will and the Concept of a Person." *Journal of Philosophy* 68.1 (1971): 5-20. Print.
- Hume, David. *A Treatise of Human Nature*, Ed. L. A. Selby-Bigge. Oxford: Oxford UP, 1978. Print.
- Mckenna, Michael. "A Hard-line Reply to Pereboom's Four-case Manipulation Argument." *Philosophy and Phenomenological Research* 77.1 (2008): 142-59. Print.
- Mele, Alfred. "Free Will and Luck: Replies to my Critics." *Philosophical Explorations* 10.2 (2007): 195-210. Print.
- Pereboom, Derk. "Hard Incompatibilism." In *Four Views on Free Will*, by John Martin Fischer, Robert Kane, Derk Pereboom, and Manuel Vargas. Oxford: Blackwell, 2007. 85-126. Print.
- Pereboom, Derk. "Defending Hard Incompatibilism." *Midwest Studies* 29 (2005): 228-47. Print.
- Ramsey, F.P. *Foundations: Essays in Philosophy, Logic, Mathematics and Economics*. Ed. D. H. Mellor. Atlantic Highlands, NJ: Humanities Press, 1978. Print.
- Wallace, R. Jay. *Responsibility and the Moral Sentiments*. Cambridge: Harvard UP, 1994. Print.