



WESTFÄLISCHE
WILHELMS-UNIVERSITÄT
MÜNSTER

› The Frankfurt-Style Cases: Philosophical Lightning Rods

John Martin Fischer



Preprints and Working
Papers of the Centre for
Advanced Study in Bioethics
Münster 2012/46



› The Frankfurt-Style Cases: Philosophical Lightning Rods

John Martin Fischer

I. Introduction

An extraordinary amount of attention has been paid to the “Frankfurt-Style Cases” (FSCs), the template for which was provided by Harry Frankfurt in his famous 1969 article, “Alternate Possibilities and Moral Responsibility.”ⁱ The prototype for the Frankfurt-style cases actually goes back to John Locke’s discussion of a man who voluntarily stays in a room which, unbeknownst to him, is locked.ⁱⁱ The Frankfurt-style cases involve a signature sort of preemptive overdetermination, and some philosophers have followed Harry Frankfurt in concluding that the examples show the Principle of Alternative Possibilities (PAP) to be false. According to (PAP), an agent is morally responsible for an action only if he was free to do otherwise.

i Harry G. Frankfurt, “Alternate Possibilities and Moral Responsibility,” *Journal of Philosophy* ; for a helpful overview of the literature, see David Widerker and Michael McKenna, eds.

ii John Locke

At this point it will be helpful to have a version of the FSCs before us:

Black is a stalwart defender of the Democratic party, despite some disappointments about Obama. He has secretly inserted a chip in Jones's brain which enables Black to monitor and control Jones's activities. Black can exercise this control through a sophisticated computer that he has programmed so that, among other things, it monitors Jones's voting behavior. If Jones were to show any inclination to vote for Romney (or, let us say, anyone other than Obama), then the computer, through the chip in Jones's brain, would intervene to assure that he actually decides to vote for Obama and does so vote. But if Jones decides on his own to vote for Obama (as Black, the old progressive would prefer), the computer does nothing but continue to monitor—without affecting—the goings-on in Jones's head.

Now suppose that Jones decides to vote for Obama on his own, just as he would have if Black had not inserted the chip in his head. It seems, upon first thinking about this case, that Jones can be held morally responsible for this choice and act of voting for Obama, although he could not have chosen otherwise and he could not have done otherwise.ⁱⁱⁱ

The contention that the FSCs establish, or help to establish the falsity of (PAP) is highly controversial. An X-phi type of philosopher might even find that most philosophers who think about such matters are not convinced that the FSCs successfully (or, at least, uncontroversially) refute (PAP). Some philosophers—including me—have contended that not only do the FSCs show (to a reasonable degree of plausibility) that (PAP) is false, but they also provide an important ingredient in a larger argument that causal determinism is compatible with moral responsibility.^{iv} More specifically, I have argued that the FSCs provide a plausibility argument for the (preliminary) conclusion that if the agent in the example is not morally responsible, it is not because he lacks freedom to do otherwise. I have further argued that causal determination in the actual sequence does not in itself and apart from ruling out alternative possibilities constitute a responsibility-undermining factor.^v Thus, I hold that the FSCs are an important part of a two-step argument to the conclusion that causal determinism is compatible with moral responsibility.

Every step of the argument is highly contentious. And not only has there been vigorous controversy about each step, there is increasing hostility among some philosophers to the continued intense scrutiny received by the FSCs. Indeed, some journals have recently decided not to publish any more articles on the FSCs, and others have explicitly (or perhaps implicitly) set a “high bar” for acceptance of such articles. This may then be a propitious occasion on which to step back from some of the spirited debates about specific features of the examples and their significance and to consider some more general “meta-questions”. In this paper I propose to do just this.

iii I originally presented such an example in John Martin Fischer, “Responsibility and Control,” *Journal of Philosophy* (1982) 79, pp. 24–40.

iv John Martin Fischer,

v John Martin Fischer,

II. The FSC Skeptic's Challenge

1. The Challenge

In developing and presenting the FSC-skeptic's challenge, I owe much to conversations with Patrick Todd. I suppose I could have called the worry, "Todd's challenge," but, although (in these conversations) Todd has crystallized the issues in a particularly sharp and helpful way, they are widely held. Here is the worry. (PAP) is deeply ingrained both in common sense and in our more reflective theorizing both in philosophy and the criminal law. It is indeed highly intuitive as well as almost universally accepted. Given this, it would take—or should take—a lot to persuade us to overturn and reject it. Further, the arguments against it seem to be "close calls"; they are highly contentious, in any case, and even if some argument or other is persuasive to some, it will not be so clearly compelling that any fair and open-minded philosopher (not already committed to a view about [PAP]) would need to accept it. Indeed, the arguments are either extremely "close calls" or are so complex that, even if one is inclined to accept them, one should not be very confident that one has really grasped all of the relevant issues. Given all of this, it just doesn't seem philosophically prudent or sensible to give up such a deeply plausible and central philosophical principle, such as (PAP).

No doubt, this challenge captures what many philosophers are inclined to think about the FSCs and their implications. It represents a legitimate set of concerns. I shall attempt to address the challenge, and I hope that in the process, we can shed some light on the basic worries that drive it, as well as what I take to be the best strategy for defending the contention that the FSCs indeed do provide an important part of an argument for compatibilism about causal determinism and moral responsibility.

I accept a broadly speaking "coherentist" methodology here, similar to what John Rawls employed in seeking what he called a "reflective equilibrium" or match between our considered judgments about cases and our general principles. I favor a version of this methodology that attempts to secure what Norman Daniels called a "wide reflective equilibrium"; on this approach, one considers a wide range of cases, both hypothetical and actual, and also a wide array of potential principles. In any case, the methodology involves sometimes re-thinking and adjusting one's considered evaluations of or "intuitions about" particular cases, and sometimes adjusting one's general principles in an effort to achieve a philosophical homeostasis. Importantly, the methodology is "holistic"; it does not simply rely on intuitions—firm and clear or otherwise—about a particular case or class of cases.

In reflecting on the FSC presented above, I think it is plausible that Black's presence and dispositions (perhaps together with other facts) makes it the case that Jones was not free to do otherwise than choose to vote for Obama and so vote. Further, it seems to me, as it seemed to Frankfurt, that Black's presence and dispositions (perhaps together with the relevant other facts) are irrelevant to Jones's moral responsibility. It would appear to follow that the fact that Jones could not have done otherwise is irrelevant to his moral responsibility. Thus, if causal determinism threatens an agent's moral responsibility, it is not in virtue of ruling out his freedom to do otherwise (his metaphysical access to alternative possibilities).

But here, as elsewhere in philosophy, things are not so simple! Various philosophers hold that it is not evident that in the FSC as developed above Jones is not free to do otherwise. And they contend that if the description of the example is filled in so that it becomes clear that Jones indeed lacks this sort of access to alternative possibilities, it becomes correspondingly *unclear* that the facts that rule out his freedom to do otherwise are indeed irrelevant to his

moral responsibility. The proponent of the claim that the FSCs prove the falsity of (PAP)—to some suitable degree of plausibility—must hold both that in the FSCs the agent is not free to do otherwise and also that he is morally responsible. And there is a tension between these two claims.

This tension is brought out nicely by what has come to be called the “Dilemma Defense.” Does the mere presence of Black, together with his dispositions and technology, really make it the case that Jones cannot at just prior to the time of his choosing to vote for Obama (say time T_2) do (i. e., choose) otherwise at T_2 ? On the first horn of the dilemma, we suppose that causal indeterminism obtains, and that it obtains specifically in the relationship between Jones’s mental states prior to T_2 and his choice at T_2 . Now it would seem that, no matter what occurs along the sequence to T_2 , Jones can still at least begin to choose to vote for someone other than Obama (or not at all) right at T_2 . So under the assumption of causal indeterminism (underwritten or made true by lack of determination in the appropriate location), Jones would seem to have at least *an* alternative possibility, truncated as it would be. Although Jones’s effort to choose otherwise would be blocked, his beginning to choose to vote for someone other than Obama (or not at all) would appear to be sufficiently robust to ground attributions of moral responsibility, on the alternative-possibilities picture of moral responsibility.

On the other horn of the dilemma, we assume a causally deterministic relationship between Jones’ prior mental states and his choice at T_2 to vote for Obama. Now the proponent of the so-called “Dilemma Defense” will contend that it would be question-begging to extract the conclusion that Jones is morally responsible for choosing at T_2 to vote for Obama. After all, the proper conceptualization of the relationship between causal determinism and moral responsibility is precisely what it at issue in the context in which the Frankfurt-style cases are invoked. So the very resources used to secure the claim that the relevant agent lacks freedom to do otherwise seems to call into question whether he can legitimately be deemed morally responsible for the behavior in question.

In previous work I have contended that the Dilemma Defense can be resisted by the proponent of the position that the FSCs effectively impugn (PAP). First, although I have not argued for this view, I have expressed the hope that certain indeterministic versions of the FSCs “work”. More specifically, I believe that the “Buffer Zone” cases developed by such philosophers as David Hunt and Derk Pereboom hold out the hope that indeterministic FSCs can be invoked to defeat (PAP).^{vi} It may indeed be the case that other indeterministic versions of the FSCs can be shown (to an acceptable degree of plausibility) to be cases in which an agent cannot do otherwise and also is morally responsible for the relevant behavior.^{vii} I find it plausible that an indeterministic Frankfurt-style case can be constructed in which the *only* remaining alternative possibilities are “mere flickers of freedom” and thus not sufficiently robust to ground attributions of moral responsibility.

To explain. I have argued (in various places) that *if* one accepts an alternative-possibilities model of moral responsibility (i. e., an approach on which freedom to do otherwise is required for moral responsibility), then one should require that the alternative possibilities in question be *robust*. After all, if one holds that alternative possibilities are the ground of moral responsibility attributions, they should not indicate the mere possibility of something else occurring; they

vi David Hunt; Derk Pereboom

vii Mele/Robb, others?

must be, as I have insisted, sufficiently *robust*.^{viii} Now I claim, without any argument, that it is at least plausible that an indeterministic FSC can be constructed in which the only alternative possibilities that remain are not sufficiently robust to ground attributions of moral responsibility; they are, as I would put it, *mere flickers of freedom*.

I would make exactly the same claim about the deterministic horn of the dilemma. I have argued that we can construct causally deterministic versions of the FSCs in which the only remaining alternative possibilities are mere flickers of freedom. Further, I have argued that these versions of the FSCs need not beg the question against the incompatibilist, simply in virtue of positing causal determinism.^{ix} The point is that, having presented a causally deterministic version of the FSC, I do not ask the reader precipitously and straightaway to judge that the relevant agent (say, Jones) is morally responsible for the behavior in question. Rather, at this stage I simply invite the reader to come to a preliminary conditional conclusion: *if* the agent (say Jones) is not morally responsible for his actual behavior, then this is not in virtue of his lacking freedom to do otherwise. That is, on this more patient and judicious approach, we first get to a preliminary conclusion that does *not* beg the question against an incompatibilist about causal determinism and moral responsibility: that it is not the *mere fact* of lack of alternative possibilities that rules out the agent's moral responsibility, if it is indeed ruled out. Then I would proceed to a second step in which I would seek to argue that causal determination in the actual sequence does not *in itself and apart from ruling out alternative possibilities* threaten moral responsibility.^x

It will be useful to have before us my elaboration of a causally deterministic version of an FSC. We assume that the story is as in the original presentation of the Frankfurt case above, except that we make it explicit that causal determinism obtains and we also are explicitly *agnostic* about the relationship between causal determinism and freedom to do otherwise. That is, we make no assumption about whether causal determinism rules out alternative possibilities.

To understand how the presence of Black under such assumptions is supposed to rule out freedom to do otherwise, assume first that Jones chooses at *T2* to vote for Obama at *T3*. Assume also that Jones exhibits a reliable sign at *T1* that indicates how he will choose to vote at *T2* (and vote at *T3*). Let's say that if he raises his left eyebrow at *T1* he will choose at *T2* to vote for the Democrat, and if he raises his right eyebrow at *T2* he will choose at *T2* to vote for the Republican. Here is how the example is supposed to work to get to the conclusion that Jones cannot at or just prior to *T2* choose otherwise:

Black checks and sees the 'prior sign' at *T1* that is associated with a subsequent vote for the Democrat—say, the [raised] *left* brow. Given that Black knows that causal determinism obtains, he can now relax, as it were; under these circumstances, Black knows that Jones in fact will subsequently choose to vote for Obama and carry out that choice. It is *also* true, given Black's device and dispositions, that if Jones were to show the sign at *T1* associated with voting for a Republican at *T2* (appropriately enough, a raised *right* brow), Black's device would swing into action and stimulate

viii I myself have not sought to give an analysis of the notion of robustness at issue. Rather, I have posited that it is at least a necessary condition of the alternative possibilities' being robust (in the required sense) that they involve voluntary behavior. For a helpful account of robustness, see Derk Pereboom:

ix John Martin Fischer, Ethics, Frankfurt-Style Compatibilism, Resp and the Actual Sequence, etc.

x John Martin Fischer, MOFW, FSC, "Compatibilism" (in Four Views)

Jones's brain so as to ensure that he chooses at $T2$ to vote for Obama and does so vote at $T3$. I claim that this additional fact, when added to the assumption of causal determinism and the fact that Black can thus be sure that Jones's showing the prior sign at $T1$ will in fact be followed by his choosing accordingly at $T2$, renders it true that Jones cannot at $T2$ choose to vote for [Romney] (or subsequently vote for Romney). These two facts *together* make it the case that Jones cannot at $T2$ choose to vote for [Romney] or carry out such a choice.^{xi}

Now note here that in this version of the case the agent (Jones) can indeed show the sign at $T1$ associated with voting for a Republican at $T2$ —he could (involuntarily) raise his right eyebrow. But, although this alternative possibility remains in place, it is obviously a mere flicker of freedom and not sufficiently robust to ground attributions of moral responsibility. So the causally deterministic version of the FSC (on the assumption of agnosticism about the relationship between causal determinism and freedom to do otherwise) is similar in this respect to the causally indeterministic versions (although I did not present a detailed analysis of the indeterministic versions): whatever alternative possibilities that remain are mere flickers of freedom and thus insufficient to ground attributions of moral responsibility. Thus, (PAP)—interpreted suitably—is shown to be false.

Return, now, to the incompatibilist's challenge above. It should now have some additional force, since it will have been seen just how much of a "close-call" the argumentation on behalf of the claim that the FSCs impugn (PAP) might appear to be. Note that the proponent of this view about the implications of the FSCs has not been able to provide an example in which there is literally *no* alternative possibility and the agent is still plausibly thought to be morally responsible for the behavior in question. Given that the argumentation is so complex and the issues so delicate, how can we be asked to give up such an extremely attractive and important principle as (PAP)? Doesn't philosophical prudence and good sense require us to proceed more cautiously here?

2. A Reply to the FSC Skeptics

As I wrote above, my favored methodology seeks a "wide reflective equilibrium," and I do not rest my case (nor do I think the case should be rested) on intuitions about the FSCs. I take a more holistic approach, in which there are various components to the argument. First, of course, I find it intuitively plausible that in the various "preferred" versions of the FSCs (the Buffer Zone cases, in particular, and my causally deterministic prior-sign case) the agent lacks the relevant kind of alternative possibilities. Further, I find it intuitively plausible that the agents' lack of the relevant kind of alternative possibility would not preclude their moral responsibility; that is, I find it plausible that if the agents are not morally responsible, this is not in virtue of lacking alternative possibilities. This is not nothing, but I do not stop here.

Notice also that there is a plausible "principle" or more general idea that explains the intuitions: moral responsibility for behavior depends on the features of the actual sequence that issues, and not on whether the agent has access to alternative possibilities. This kind of principle can both provide support for the intuitions in question and also (together with those

xi Fischer, "The Frankfurt Cases: The Moral of the Stories," p. 326.

intuitions) motivate a “paradigm shift” in our thinking about moral responsibility. Whereas traditionally we were concerned with whether the agent could have done otherwise, we now scrutinize the features of the actual sequence leading to the action with an eye to ascertaining whether the actual sequence contains any responsibility-undermining features in themselves and apart from ruling out alternative possibilities.

But there is more even than intuitions in equilibrium with general principles. The “actual-sequence” theory of moral responsibility has an Error Theory that explains the traditional appeal of (PAP). We can begin by using the terminology of Harry Frankfurt, who distinguishes two kinds of freedom: freedom to do otherwise and acting freely. Typically, it would be assumed that whenever an agent acts freely, it is also the case that he is free to do otherwise. The FSCs, however, with their distinctive structure involving preemptive overdetermination, can help us analytically to prize apart the two kinds of freedom and to see that it is not necessary that they are co-instantiated. More specifically, the FSCs appear to be cases in which an agent acts freely, but is not free to do otherwise. Perhaps people mistakenly assume that the two kinds of freedom necessarily go hand in hand, and thus they do not distinguish the two kinds of freedom for the purposes of attributions of moral responsibility. The Error Theory has it that either people fail analytically to prize apart the two kinds of freedom, or having done so, they mistakenly assume that they necessarily go together, so that in linking moral responsibility with freedom to do otherwise, one is also (necessarily) linking moral responsibility with acting freely. This is an error, because the FSCs show that an agent can be morally responsible in virtue of acting freely, even in the absence of freedom to do otherwise. It is an error that is easy enough to make, unless one thinks carefully about examples with the signature structure of preemptive overdetermination originally sketched by John Locke and developed further by Harry Frankfurt and his followers. And because the examples have been developed with considerably greater sophistication in recent years, it is an error that can be identified and highlighted more effectively nowadays.

The Error Theory can of course also be elaborated with my favored distinction between two kinds of control: regulative and guidance control. Whereas regulative control involves access to alternative possibilities (freedom to do otherwise), guidance control need not. It is perhaps natural to assume that the two kinds of control always (and necessarily) go together, i. e., that when an individual has guidance control, he also must have regulative control. So it would be natural just to focus on one of these distinct kinds of control, if one analytically separates them in the first place. Given the salience of regulative control, it would be a natural enough error to make simply to link moral responsibility to regulative control. But the FSCs show that this is indeed an error; they show that there can be cases in which an agent is morally responsible in virtue of exhibiting guidance control, even in the absence of regulative control.

So there is a package of considerations that are part of the holistic evaluation of the FSCs, on the methodology of seeking a wide reflective equilibrium: intuitions explained by a general principle, together with an Error Theory that explains the appeal of an alternative principle. I contend that there is even more. It is desirable to have a “picture” that corresponds to one’s theory of moral responsibility or perhaps a deep explanation of the “value” of moral responsibility. And in previous work I have sought to provide exactly this. On my approach, the value of moral responsibility is the value of a certain distinctive kind of self-expression, rather than the value of “making a difference” (which is associated with the traditional view that freedom to do otherwise [regulative control] is required for moral responsibility). On my view, the value of moral responsibility (i. e., the value attached to acting in such a way as to be deemed morally responsible) is the value of making a statement, not making a difference.^{xiii}

That there is a plausible foundational “picture” associated with the actual-sequence package of views further enhances its plausibility. Is this web of interlocking and mutually supporting components enough to overthrow the traditional adherence to (PAP) and the alternative-possibilities model of moral responsibility? I am not sure, but I would suggest that the actual-sequence theory at least has a fighting chance here; and it should be evident that its appeal do *not* rest solely on intuitions about cases (difficult and delicate as they are). But I wish to point out that there is even more to be said on behalf of this approach, given the holistic methodology I have adopted. It is important to keep in mind that a shift from the alternative-possibilities to the actual-sequence approach allows us to side-step traditional worries (say) about the relationship between causal determinism (or God’s foreknowledge) and moral responsibility. Thus such a shift opens the possibility of making genuine philosophical progress with respect to the issue of the compatibility of causal determinism (or God’s foreknowledge) and moral responsibility.^{xiii} This is another big advantage of an actual-sequence view, in the context of a thoroughly holistic philosophical cost-benefit analysis. All in all, I wouldn’t bet against an actual-sequence model of moral responsibility.

III. FSCs and Dialectical Stalemates

I have no illusions that the reflections in the previous section will move all skeptics all the way (or even much of the way) to an actual-sequence approach to moral responsibility. A residual worry is as follows. None of the examples I have presented (or sketched) above—the deterministic and indeterministic FSCs I have claimed are the most promising—contain *no* alternative possibilities. There are, after all, those pesky “flickers of freedom” that I claim are *mere* flickers of freedom (and thus not sufficiently robust to ground attributions of moral responsibility). But flickers they are—and some will wish to fan the flickers of freedom to get what Ronald Reagan used to call a “prarie fire of freedom”. In any case, some skeptics will remain unconvinced precisely because no one has yet produced an FSC in which it is uncontroversially true that there the agent is morally responsible for his behavior and he has *no* alternative possibilities. The examples, then, *in themselves*, do not seem to do the trick; they always appear to fall short of what is desired (and perhaps needed).

I wish to point out, however, that it is not necessary, and arguably not reasonable to expect, that the examples themselves will “go all the way” to refuting (PAP). One of course might see how far we could get in seeking to construct examples in which the agent is uncontroversially morally responsible for his behavior and there are *no* alternative possibilities. But it is not a philosophical disaster (for the critics of [PAP]) if we do not fully succeed. This seems to me to be another philosophical context that can be illuminated by reference to the dialectical structure I have dubbed, “Dialectical Stalemates”.^{xiv} Here is a description of these “black holes in dialectical space/time”:

Frequently in philosophy we are engaged in considering a certain argument (or family of arguments) for some claim *C*. The argument employs a principle *P*. Allegedly, *P* supports *C*. Now the proponent of the argument might be called upon to support the principle, and he may do so by invoking a set of examples (or other considerations). Based on these examples (or

xii Fischer, Self-expression, etc.

xiii John Martin Fischer, Frankfurt-Style Compatibilism”

xiv John Martin Fischer, MOFW, “Epicureanism About Death and Immortality”

other considerations), he argues that the principle and thus also the philosophical claim are to be accepted.

But the opponent of the argument may respond as follows. The examples are not sufficient to establish the principle *P*. One could embrace all the examples and not adduce *P* to explain them; rather, it is alleged that a weaker principle, *P'*, is all that is decisively established by the examples (or other considerations). Further, *P'*, in contrast to *P*, does not support *C*. Finally, it is very hard to see how one could decisively establish *P*. One reason it is so difficult is that it at least appears that one cannot invoke a particular example which would *decisively* establish *P* without begging the question in a straightforward fashion against either the opponent of *P* or the opponent of *C*. Further, it *also* seems that one cannot invoke a particular example which would *decisively* refute *P* without begging the question against the proponent of *P* or the proponent of *C*. These conditions make out a distinctive—and particularly precarious—spot in dialectical space.

I shall call contexts with roughly the above form, “Dialectical Stalemates.”^{xv}

Of course, although in the above passage I wrote in terms of seeking to defend a principle *P*, exactly the same considerations will apply to an attempt to refute a principle *P* (such as PAP). Elsewhere I have argued that the debates about the Consequence Argument, and, in particular, its crucial ingredients—the Principle of the Fixity of the Past and the Principle of the Fixity of the Natural Laws—have reached Dialectical Stalemates.^{xvi} Additionally, I have argued that the debates about the badness of death for the individual who has died have also reached a Dialectical Stalemate.^{xvii} Here I claim that the debates about FSCs also have all of the dialectical stigmata of these notoriously intractable philosophical standoffs, but that they can be illuminated by the analytical framework that I have sketched previously.

It will be helpful to start with the examples offered by Thomas Nagel and others that purport to show that death (thought of as an experiential blank) can indeed be bad for the individual who has died. Nagel starts with a case in which a person is betrayed behind his back, although he never finds out about it and, we can stipulate, never has any unpleasant experiences (or is deprived of pleasant experiences) because of these regular betrayals.^{xviii} Nagel claims that these betrayals are bad for the individual, even though they never have a negative impact on his experiences.

It is interesting to note Nussbaum’s reply:

Nagel does not make it clear exactly how an event located completely outside a life’s temporal span diminishes the life itself. The cases he actually analyzes are not by themselves sufficient to show this, since in each of them a subject persists, during the time of the bad event, who has at least a strong claim to be identical with the subject to whom the bad event is a misfortune. In the betrayal case, this subject is clearly the very same, and is a subject of possible, if not actual, experience in relation to that event. ... Where death is concerned, however, there is no subject at all on the scene, and no continuant. So it remains unclear exactly how the life that has ended is diminished by the event.^{xix}

xv Fischer, MOFW, p. 83.

xvi Fischer, MOFW, pp. 67–86.

xvii Fischer, “Epicureanism About Death and Immortality,” esp. pp. 112–118.

xviii Thomas Nagel, “Death”.

xix Martha Nussbaum, *The Therapy of Desire* (Princeton: Princeton University Press, 1994), pp. 205–6.

The example can be (and has been) made more sophisticated.^{xx} And similar examples have been proposed in the literature on the metaphysics of death. But the challenge is that, however we adjust the examples and make them more refined, there will still be a subject present; and thus the examples will, strictly speaking, fall short of their proponents' goal: to present an example that in itself *entails* that death (conceived of as nonexistence) can be a bad thing for the individual. And if we were to propose a case in which the subject has indeed gone out of existence and yet we claim that intuitively some condition or event is a bad thing for that individual, this will of course be unhelpful, given the dialectical situation: to propose a case in which the individual has gone out of existence is too close to proposing a case of death being a bad thing for an individual—which is precisely the question at issue. The parallel with the dialectical context of the FSCs should be evident: the more we add ingredients to rule out (say) Jones's freedom to do otherwise, the closer we come to assuming causal determinism; and this can seem to render the arguments of the proponents of the FSCs (at least construed in a simple, one-step manner) question-begging.

I do indeed think we can learn from the dialectical context involving the metaphysics of death. In a Dialectical Stalemate, no example *can in itself decisively establish* the relevant conclusion (without begging the question). Nussbaum is correct to point out that none of the examples proposed in the literature on death's putative badness can in themselves decisively establish that death can be a bad thing for the individual who dies. But it is important to keep in mind that we have various options when faced with a Dialectical Stalemate:

I do not however think that Dialectical Stalemates should issue in philosophical despair. An opponent of the principle under consideration may demand that its proponent provide examples which absolutely *require* one to accept the principle. But I would claim that this is *unreasonable*. It may even be true that it is *necessarily* the case that if a philosopher argues for a certain general principle by giving examples, a weaker principle can be found that is the strongest principle the examples support (strictly speaking). The crucial issue becomes whether it is plausible to accept the stronger principle, if one accepts the weaker principle...^{xxi}

I believe that, although we are not *required* to conclude from Nagel's example (and related examples) that death can be bad for the individual who dies, it is nevertheless *reasonable* to come to this conclusion, given that we hold that it is indeed bad for an individual to be betrayed behind his back (in the scenario described by Nagel, perhaps embellished in certain ways). That is, I hold that it is plausible that *if* the individual is harmed in Nagel's case, he is *also* harmed by death. That is, I don't think that the difference between Nagel's case and the case to which he seeks to extrapolate the result—the case of death—*makes a difference* as regards the crucial issue of the badness of the relevant feature. Nagel's example (and related examples) are artfully chosen so that it is plausible to extrapolate from them to the admittedly different case of death; of course, if there were no difference, there would be no need for extrapolation—but then again one could not maintain the judgment of badness without (presumably) begging the question.

xx Fischer (White example)

xxi Fischer, MOFW, p. 85.

I would say exactly the same thing about the dialectic involving the FSCs. Although we are not *required* by the various FSCs in themselves to jettison (PAP), I contend that it is nevertheless *reasonable* to do so, given that we hold that the agents in the FSCs are indeed morally responsible. The FSCs in question are cases in which it is plausible that an agent is morally responsible, even though he has only mere flickers of freedom available to him; thus, they strongly suggest that an agent would be morally responsible, even if he had *no* alternative possibilities at all. The FSCs are artfully constructed to support the plausibility of this extrapolation to the admittedly different case in which there are literally *no* alternative possibilities at all. Just as with the proposals concerning the badness of death, in the case of the FSCs we have good reason to reject the relevant principle: (PAP). Just as with the examples related to death, the difference between the proposed cases and the scenario to which we seek to extrapolate the result *does not seem to make a difference*.

I conclude that at least some of the residual anxiety about rejecting (PAP) that stems from the fact that no existing FSC is uncontroversial example in which an agent is morally responsible and has no alternative possibilities is misplaced; it is based on unreasonable philosophical expectations. Of course, if it could be argued that in every proposed FSC there is at least one robust alternative possibility, then this would be quite a different matter. But this is where the argument would need to focus, and thus far I have not been convinced by any such argument.