Westfälische
Wilhelms-Universität
Münster

The Irrelevance of Counterfactual Intervention

John Martin Fischer

45

Preprints and Working
Papers of the Centre for
Advanced Study in Bioethics
Münster 2012/45

wissen.leben
WWU Münster

NORMENBEGRÜNDUNG
IN MEDIZINETHIK
UND BIOPOLITIK
KOLLEG-FORSCHERGRUPPE | WWU MÜNSTER

› The Irrelevance of Counterfactual Intervention

John Martin Fischer

I. Introduction

I have argued for an "actual-sequence" model of moral responsibility. On an actual-sequence approach, moral responsibility does not require the kind of control that involves "freedom to do otherwise" or, in a more regimented manner of expression, genuine access to alternative possibilities. If an actual-sequence model of moral responsibility could be defended successfully, this would bring tremendous philosophical benefits. In particular, we would be able to side-step the traditionally contentious and apparently intractable debates about whether human beings can be free to do otherwise, if there are propositions that were true in the past that specify how we will behave in the future, if God (understood in a certain way) believed in the past that we would behave as we do, or if causal determinism obtains. These debates lead to what I have dubbed, "Dialectical Stalemates"—black holes in philosophical space-time in which progress is hard to come by. If we do not need freedom to do otherwise in order to be morally responsible, it is possible that we can indeed make philosophical progress on moral responsibility. Switching from the traditional "alternative-possibilities" model to an actual-sequence model is nothing less than a paradigm shift. There is a lot at stake, then, in evaluating the contention that we ought to adopt an actual-sequence approach to moral responsibility.

There are various different ways of elaborating and defending an actual-sequence approach to moral responsibility. As far as I am concerned, I welcome these alternative approaches; this is one context (the context of defending an actual-sequence approach) in which, perhaps ironically, I embrace alternative possibilities! Alternative strategies for defending the actual-sequence approach to moral responsibility include what might be called "Strawson-style" and "Dennett-style" argumentation. I have explored a "Frankfurt-style" approach to the defense of the actual-sequence model of moral responsibility. I wish here to elaborate the defense a bit further in light of some objections, and I shall highlight the way in which an actual-sequence approach to moral responsibility can help to address challenges to our responsibility stemming from both causal determinism and causal indeterminism.

II. The Frankfurt-Style Cases

In order to "situate" my strategy for addressing the luck problem within a larger context, it might be helpful to recall the so-called "Frankfurt-style" counterexamples to the Principle of Alternative Possibilities (PAP). These much-discussed examples were originally introduced in contemporary philosophy by Harry Frankfurt in order to impugn (PAP), according to which moral responsibility requires the kind of control that involves freedom to choose and do and otherwise.[i] Here is an updated version of a Frankfurt example:

> 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.[ii]

It seems to me that Black's presence (as described in the example), perhaps together with other features, makes it the case that Jones cannot choose or do other than he actually does. Further, it seems to me that Black's presence (in the context of those other features) is *irrelevant* to Jones's moral responsibility. It might be helpful to have before us Frankfurt's statements on behalf of the contention that Black's presence is irrelevant to Jones's moral responsibility:

> The fact that a person could not have avoided doing something is a sufficient condition of his having done it. But, as some of my examples show, this fact may play no role whatever in the explanation of why he did it. It may not figure at all among the circumstances that actually brought it about that he did what he did, so that his action is to be accounted for on another basis entirely. … Now if someone had no alternative to performing a certain action but did not perform it because he was unable to do otherwise, then he would have performed exactly the same action even if he *could* have done otherwise. The circumstances that made it impossible for him to do otherwise could have been subtracted from the situation without affecting what happened or why it happened in any way. Whatever it was that actually led the person to do what he did, or that made him do it, would have led him to do it or made him

i Frankfurt, "Alternate Possibilities and Moral Responsibility." For a selection of papers on Frankfurt-style examples, see David Widerker and Michael McKenna (eds), *Moral Responsibility and Alternative Possibilities: Essays on the Importance of Alternative Possibilities* (Aldershot, U. K.: Ashgate, 2003).

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

> do it even if it had been possible for him to do something else instead. … When a fact is in this way irrelevant to the problem of accounting for a person's action it seems quite gratuitous to assign it any weight in the assessment of his moral responsibility.[iii]

I agree with Frankfurt's intuition that it would be "quite gratuitous" to assign any weight to Black's presence in assessing Jones's moral responsibility. After all, Black's device, although present, is *untriggered*. I think that the Frankfurt-style examples help to provide motivation for an "actual-sequence" approach to moral responsibility, according to which moral responsibility attributions depend on (possibly dispositional or modal) features of the actual sequence, rather than on the availability of genuinely open alternative possibilities. The mere presence of certain sorts of *untriggered ensurers* (such as Black's device) rules out alternative possibilities without in any way affecting the relevant features of the actual sequence.

I like to call Frankfurt's intuition here, "Frankfurt's Quite Gratuitous Point". But I concede that the name suffers from a kind of infelicitous, if delicious, ambiguity. I shall follow David Palmer in crystallizing a principle—The Irrelevance Principle (IP)—that arguably captures Frankfurt's point here:

(IP) If a fact is irrelevant to a correct account of the causal explanation of the person's action, then this fact is irrelevant to the issue of the person's moral responsibility.[iv]

III. The Irrelevance Principle

It seems to me that the Irrelevance Principle, suitably interpreted, is true. But various philosophers have found fault with it (or a closely related principle seeking to capture the basic idea). For example, David Widerker has presented some interesting proposed counterexamples to (IP).[v] In a recent paper, David Palmer has presented a variant of one of Wideker's examples (which Palmer deems a "clear counterexample"):

> Suppose that Jones decides to break a promise for some personal gain despite knowing that it is morally wrong for him to do this. He does not decide to break his promise *for* the reason that it is morally wrong. He makes the decision *in spite* of knowing this, deciding to break his promise simply for self-interested motives. In other words, since Jones decides to break his promise for personal gain, it is true of him that he would have made the same decision whether or not he believed his decision to be immoral.[vi]

Palmer goes on to analyze the example as follows:

> The fact that Jones knows that his decision to break his promise is morally wrong is irrelevant to a correct account of the causal explanation of what he did and so, according to (IP), it would be irrelevant to the issue of his moral responsibility. But surely

iii Frankfurt, "Alternate Possibilities and Moral Responsibility", pp. 836–7.

iv David Palmer, "Deterministic Frankfurt Cases," unpublished manuscript, University of Tennesee, Knoxville, Department of Philosophy.

v David Widerker, "Frankfurt's Attach on the Principle of Alternative Possibilities: A Further Look," Philosophical Perspectives 14 (2000), pp. 181–201.

vi Palmer, "Deterministic Frankfurt Cases". ms. p. 13.

the fact that Jones knows it would be morally wrong to break the promise and yet does so anyway *does* bear on his moral responsibility. A person is more blameworthy, it would seem, if he knows his action is morally wrong and acts anyway in spite of knowing this, than if he does not know it to be morally wrong.[vii]

Here is a pair of examples that also seems to impugn (IP), "No-Sharks" and "Sharks":

No-Sharks: John is walking along a beach, and he sees a child struggling in the water. John believes that he could save the child with very little effort (which we suppose is true in the example), but he is disinclined to expend any energy to help anyone else. He decides not to try to save the child, and he continues to walk along the beach. Unfortunately, the child drowns.

Sharks: John is walking along a beach, and he sees a child struggling in the water. John believes that he could save the child with very little effort, but he is disinclined to expend any energy to help anyone else. He decides not to try to save the child, and he continues to walk along the beach. Unbeknownst to John, if he had jumped in, a patrol of sharks that infested the water between the beach and the struggling child would have eaten him.[viii]

It seems that John is morally responsible for failing to save the child in No-Sharks, but not in Sharks, and yet the presence and dispositions of the sharks—the only differentiating factor—plays no role in the explanation of John's behavior in the actual sequence (behavior that arguably constitutes his failure to save the child). It is admittedly somewhat difficult to specify what exactly explains a failure or omission; in any case, it is at least plausible that the presence of the Sharks does not explain John's failing to save the child. The Sharks certainly are not a causal link in the actual sequence leading to John's behavior—behavior that constituted, in the circumstances, his failing to save the child.[ix]

Admittedly, cases involving omissions present delicate issues; here is a pair of cases (due to Carolina Sartorio) involving moral responsibility for consequences:

SWITCH: A train is running out of control down a track. The train approaches a switch. Flipping the switch would send the train down a side track for a while, but the tracks reconverge up ahead, before the location of a victim who has been tied to the track. An agent, who is standing by the switch, has reason to believe that a large fragment of track is missing from the main line between the switch and the victim's location. So he has reason to believe that, unless he flips the switch, the train will continue on the main track and derail, and thus the victim will survive. Given that he wants the victim to die, he flips the switch. As a result, the train travels on the side track, then back on the main track again after the tracks come back together, and ends up killing the victim.

vii Palmer, "Deterministic Frankfurt Cases," ms. p. 13.

viii "Sharks" was first introduced in John Martin Fischer and Mark Ravizza, Responsibility and Control: A Theory of Moral Responsibility (New York: Cambridge University Press, 19980, p. 125.

ix Here it might be helpful to distinguish the explanation of John's actual failure to save the child from the explanation of the obtaining of the fact that, no matter what John tried, he would have failed to save the child. This is, as it were, a "modalized omission," or perhaps a modalized fact involving an omission. For discussion of modalized facts in the context of moral responsibility for omissions, see Fischer and Ravizza, Responsibility for Consequences, pp. 102–3.

Distinguish two scenarios of this kind:

SWITCH 1 A fragment of the main track between the switch and the victim's location was indeed missing.

SWITCH 2 The main track was intact (unbeknownst to the agent, someone had reconnected it earlier.[x]

Sartorio goes on to analyze these cases as follows:

> The victim dies in both cases. Now, it seems that whereas the agent is responsible for the victim's death in SWITCH 1, he is not responsible for the death in SWITCH 2. In SWITCH 2, the fact that the victim would have died even if he had failed to flip the switch seems to relieve the agent of responsibility for the death (he is still responsible for acting with a bad intention, but not for the victim's death). … This suggests that, when an agent is responsible for an outcome, factors that are not links in the causal sequence issuing in the outcome can ground the agent's responsibility for the outcome.[xi]

Despite the proposed counterexamples, I continue to believe that (IP) captures a very important—indeed, a crucial—insight. Further, I am inclined to believe that the counterexamples all exploit a vagueness in (IP), and that once this vagueness is removed in favor of a certain more specific interpretation, (IP) will emerge unscathed. Recall the formulation above:

(IP) If a fact is irrelevant to a correct account of the causal explanation of the person's action, then this fact is irrelevant to the issue of the person's moral responsibility.

The problematic vagueness is in the phrase, "the issue of the person's moral responsibility." Because actually there are various different issues pertaining to a person's moral responsibility. I would distinguish at least three such issues (or kinds of issues.) There is the issue of the *degree* of one's moral responsibility (or perhaps the degree of one's praiseworthiness or blameworthiness). There is also the issue of the *content* of one's moral responsibility, that is, the specific items for which one is morally responsible. Here we might include moral responsibility for actions, omissions, consequences (construed as particulars or universals), emotional reactions, traits of character, and so forth. And obviously we may have variation even within each category. So, for example, if someone is morally responsible for an action, it still remains to specify for exactly *which* action; similarly, if an agent is morally responsible for a consequence, it remains to specify for exactly *which* consequence, and so forth. Finally, there is the question of whether an agent (in a given context) is morally responsible *to any degree for at least something* (in the context). We might call this issue the question of whether the agent is "morally responsible at all" in the context.

Note that Widerker's example (as modified by Palmer), the case of Jones deciding to break a promise, pertains to the first question—the question of degree. But, as far as I can tell, it does not imply anything problematic about the question of moral responsibility at all; Jones is clearly morally responsible, and even more so, given his knowledge of the moral wrongness of his behavior. Similarly, both the pairs of cases involving omissions (No-Sharks and Sharks) and

x Carolina Sartorio, "Actuality and Responsibility," Mind 120 (2011), pp. 1080–81.

xi Sartorio, "Actuality and Responsibility," pp. 181–2.

the pairs of cases involving consequences (Switch1 and Switch2) pertain to the second question—the question of content. But, as far as I can tell, they do not imply anything problematic about the question of "moral responsibility at all".

I would thus suggest that the proper interpretation of (IP) is:

(IP) If a fact is irrelevant to a correct account of the causal explanation of the person's action in a given context, then this fact is irrelevant to the issue of the person's "moral responsibility at all," i. e., to the issue of whether the agent is morally responsible *to at least some degree for at least something in the context in question*.

On this interpretation of (IP), none of the purported counterexamples is successful. I believe this interpretation captures the very powerful idea Frankfurt put his finger on. And it is interesting to see that, not only does the Irrelevance Principle help to support the compatibility of causal determinism and moral responsibility, but for an importantly similar reason it supports the compatibility of causal *indeterminism* and moral responsibility. I turn to this point in the next section.

IV. Randomizing the Frankfurt-Style Cases

In the Frankfurt-style case discussed above, Black is a "counterfactual intervener", to use a term I coined many years ago. More specifically, he is an "untriggered ensurer" of the relevant result—that Jones votes for Obama. Now imagine that, although Black is still progressive, he has a wild streak (or perhaps a mischievous desire not to be too predictable). He thus decides to flip a fair coin (which, we assume here, is a genuinely random process [whatever that requires]) just prior to Jones's decision in the voting booth. Black stipulates that if the coin comes up heads, he will simply allow Jones to vote for Obama (given that Black sees that Jones is about to vote for Obama). He further stipulates that if the coin comes up tails, he will trigger an intervention in Jones's brain sufficient to cause Jones to vote for Romney instead (despite Black's seeing that Jones is about to vote for Obama—absent intervention). Let us suppose that things actually develop along the causal sequence leading to Jones's decision to vote for Obama and his voting for Obama exactly the way they do in the original case; that is, Jones votes for Obama on his own. So we are here supposing that the coin comes up heads just prior to Jones's decision and thus that Black does not trigger a pre-emptive intervention on behalf of those pesky Republicans.

Note that in the randomized Frankfurt-style case the relationship between Jones's prior mental states and his choice to vote for Obama is indeterministic; the prior states of Jones's mind do not entail Jones's choice. Some skeptics about the possibility of control of the kind that underwrites moral responsibility in indeterministic contexts have pushed the view that when the prior mental states of the agent (together with background conditions and the laws of nature) do not entail the agent's choice, then his choice is not really up to the agent in the requisite way (the way required for moral responsibility). They have taken it that when the entailment fails, the choice is arbitrary or a matter of luck; in any case, according to those who worry about the "luck problem", insofar as the entailment fails, the choice cannot be said to be attributable to the agent in the way required for moral responsibility.

Some such skeptics invoke the "Rollback Argument," according to which we could (in our imaginations) roll back the universe to (say) the beginning of an agent's deliberations and let it go forward. In some such "replays" the agent will indeed choose as he actually chooses, but in some replays he will not. Given this fact, it can seem as if the agent's actual choice is not really up to him in the sense required for moral responsibility. After all, everything internal to his

mind—all his relevant mental states—are held fixed, and yet in some replays he decides and acts differently from the way in which he actually decides and acts.

Others simply invoke the fact of Bare Transworld Differences to explain their conclusion that it is not really up to Jones how he decides (and acts). Given the lack of entailment between the agent's prior mental states (together with background conditions and laws of nature) and his choice, there will be possible worlds in which everything about the agent's prior mental states are held fixed (along with the background conditions and laws of nature) and the agent decides and acts differently from how he actually decides and acts. Again, given this fact, it can seem as if the agent's actual choice is not really up to him in the sense required for moral responsibility. That is, it seems that it is not really the agent who makes the difference with respect to his decision.

But note that the Randomized Frankfurt-style case gives us an elegant way of responding to the skeptics who worry (in various ways) about the luck problem. So let's think again about the Randomized Frankfurt-style case. We can in our imaginations roll the universe back to the time of the beginning of Jones's deliberation in the voting booth and let it go forward to the time of Jones's choice and then his action. If we were to do this multiple times (again, in our imaginations), we would note that on some replays Jones does in fact vote for Obama, and on others he does not. (This follows from—among other things—the fact that Black flips a fair coin just prior to Jones's decision, and that if the coin were to come up tails, Black would trigger an intervention sufficient to cause Jones to vote for McCain.) And yet everything internal to Jones's mind—all his beliefs, preferences, acceptances of moral principles, and personality traits—are held fixed in the multiple replays.

It is important to see that the actual sequence involving Jones's deliberations and leading to his decision to vote for Obama is exactly the same in the original Frankfurt-style case and the Randomized Frankfurt-style case. Thus, if we were inclined to think that the glue that has to hold an agent's prior mental states together with his decision is present in the original Frankfurt-style case, we should think that it is *also* present in the Randomized Frankfurt-style case. In my view, although this may be contentious in some quarter, we could even think of the actual sequence in the original Frankfurt-style case as causally deterministic. Given this assumption, we could say that it should be uncontroversial that the glue is present in the original Frankfurt-style case, and given that the intrinsic features of the actual causal sequences leading to the choice are the same in the original and randomized Frankfurt-style cases, it follows that the glue is present in the Randomized Frankfurt-style cases. And yet in the latter cases the relationship between Jones's prior mental states and his choice is characterized by causal indeterminism. So the mere fact of the applicability of the Rollback Argument does not rule out the presence of the sort of control required for moral responsibility. This at least suggests that the mere fact of causal indeterminism in the relationship between an agent's prior mental states and his decision does not rule out the control required for moral responsibility.

Exactly the same sort of analysis can help to address the problem of Bare Transworld Differences. Note that, given that Black flips a fair coin in the circumstances, there are possible worlds in which everything about Jones's prior mental states is the same as in the original case and yet he decides to vote for McCain (and does so vote). These represent "bare transworld differences". And yet, as above, the glue that holds us together as morally responsible agents—the linkage between our prior mental states and our decisions that is at least necessary, if not sufficient, for moral responsibility—is present in the Randomized Frankfurt-style case (insofar as it is assumed to be present in the original Frankfurt-style case). And thus the mere fact of bare transworld differences does not rule out the presence of the sort of control required for

moral responsibility. Again, this suggests (in the absence of further worries) that the mere fact of causal indeterminism in the relationship between an agent's prior mental states and his decision does not rule out the control required for moral responsibility.

Whereas in the original Frankfurt case, Black is an untriggered ensurer, in the Randomized Frankfurt case Black is an untriggered pre-empter. In both cases Black is equally a counterfactual intervener, and, as such, is irrelevant to Jones's "moral responsibility at all". It is interesting to observe that we can address both the central problem for moral responsibility posed by causal determinism and the central problem for moral responsibility posed by causal indeterminism in a similar way: by invoking cases involving counterfactual intervention. The intuitive engine driving the arguments is the irrelevance of counterfactual intervention.

V. The Paradigm of the Pilgrimage

In literature as well as philosophy our lives are often depicted as "journeys". Perhaps Homer's *The Odyssey* was the first great book in this tradition, which has continued through James Joyce's *Ulysses* to the present day. One moral of these stories is that we should appreciate and value the journey itself; the journey itself has intrinsic value apart from achieving any goal or arriving at certain destinations. But my point is a bit different. I claim that we can understand our lives as paths that contain no alternative possibilities—no paths that branch off the main path. More specifically, we can so understand our lives consistently with their having great value and meaningfulness—as much value and meaningfulness as we already suppose them to have. And we can so interpret our lives consistently with thinking of ourselves as fully morally responsible agents.

In my view, the switch from an alternative-possibilities model of moral responsibility to an actual-sequence model is nothing short of a paradigm shift. And I have suggested that an appropriate metaphor or picture that underlies the actual-sequence approach is the pilgrimage. In a pilgrimage we may suppose that the path is entirely determined or laid out in advance, and yet the pilgrims can nevertheless achieve great and profound spiritual benefits from the journey. The idea is that the possibilities for spiritual growth and transformation do not stem from making choices among different paths or ways of getting to the destination; rather, the meaningfulness of the journey has its roots in something more basic or "direct". Even though the pilgrim's journey is entirely laid out in advance, he can still gain great value from walking this pre-determined path in a certain way—perhaps with mindfulness or an openness to the possibilities of meaning and growth.

At the very end of *The Metaphysics of Free Will: An Essay on Control*, I wrote:

> The future may—or may not—contain more than one genuinely open path: I do not know. It is quite natural to think of the future as open, but it may turn out that the various paths I picture in my mind are mere tantalizing chimeras. Employing a slightly different metaphor, there is just one line extending from the present into the past, and the future may indeed be symmetric—there may be just one line extending into the future. But even so—even if there is just one available path into the future—I may be held accountable for *how I walk down this path*.[xii]

xii John Martin Fischer, The Metaphysics of Free Will: An Essay on Control (Oxford: Blackwell Publishers, 1994), p. 216.

About this passage, Gary Watson wondered whether my invocation of a *way of taking the path* implicitly presupposes that there are more than one such way available to the agent, in which case we do not still have an "actual-sequence" approach to moral responsibility.[xiii] But I do not think that moving from talking about moral responsibility for action or behavior to talking about moral responsibility for a way of acting or behaving requires a corresponding move to alternative possibilities. Why would conceiving of the content of our moral responsibility—what we are morally responsible for—in terms of a way of taking the path of life involve a requirement of alternative possibilities any more than conceiving our the content of our moral responsibility in terms of actions? Michael Zimmerman puts this point well:

> I find that Fischer's final rendition of his 'new paradigm' would be improved if slightly reworded. To say that, even if there is just one path available into the future, I may be held accountable for how I walk down this path, suggests (to me, at least) that I have alternative ways of walking down the path open to me. One wonders, then, whether we should be talking of one path or two… What I think Fischer should have said is this: even if there is just one available path into the future and just one available way of walking down it, I may be held accountable for walking down it in that way.[xiv]

Right. And consider the famous pilgrimage from Le Puy in south-central France to the great cathedral at Santiago de Compostela in northwestern Spain. This pilgrimage was laid out in the medieval *Pilgrim's Guide*.[xv] The pilgrimage is "pre-set" or laid out in advance, and yet this doesn't diminish the potential for spiritual growth, transformation, and meaningfulness for the pilgrims. If there are alternative paths along the way, or possible "short-cuts", I do not see how their availability in any way matters for the meaningfulness of the pilgrims taking this journey; they just seem to be irrelevant. Similarly, whether the pilgrim has a choice of when exactly to start each morning or even what to focus on as she travels the route does not seem to be relevant to the fact that taking the journey can be deeply and profoundly meaningful; again, these alternative possibilities, although undeniably present, do not seem relevant to the issue of meaningfulness. I do not even think that the possibilities for spiritual growth, transformation, and meaningfulness in a pilgrimage hinge on whether the pilgrim had alternative possibilities in choosing to take the journey in the first place.

Consider, also, the great Muslim pilgrimage or "hajj" that culminates in the holy city of Makkah in Saudi Arabia. Recently, the Chinese have built a railway that can transport pilgrims in the last phase of the journey. So the pilgrims have a choice about how to take part of the pilgrimage, but does this really matter to the meaningfulness of the journey? That is, if we sup-

xiii In "Some Worries About Semi-Compatibilism," Journal of Social Philosophy 29 (1998), p. 137, Gary Watson writes:
This affirmation, incompatibilists might complain, is a rhetorical flourish to which Fischer is not strictly entitled. Here the path metaphor seems a bit misused. In the abstract sense required by the argument, a 'way' is of course a path, a metapath, perhaps, of which there is only one if determinism is true. The aspiration to define one's own way … might be called an ideal of autonomy. Can we understand this ideal without presupposing ... alternative possibilities?

xiv Michael Zimmerman, "Book Review of John Martin Fischer's, The Metaphysics of Free Will:An Essay on Control," Canadian Philosophical Reviews 16 (1996), p. 344.

xv The Pilgrim's Guide to Santiago de Campostela, A. Gazetteer, ed and trans. Annie Shaver-Crandell and Paul Gerson, with the assistance of Alison Stones (London: Harvey Miller, 1995).

pose that a Muslim pilgrim decides to take the traditional route, rather than the train, does this really enhance the potential for meaningfulness in completing the journey to Makkah? Surely the meaningfulness of the experience for the pilgrim does not even in part depend on pride in making an appropriate choice between alternative routes at the end.

There is a famous rock-climbing ascent to the top of El Capitain in Yosemite. There is only one route up to the very top. The topological map of the route spels out exactly what a climber is supposed to do at each stage (e. g., climb X feet, stope, climb Y feet, stop, move left into a different crack, stop, etc.) People who accomplish this climb—perhaps the most famous in the world—are justly proud of their accomplishment, although the final ascent is entirely pre-set and without alternative possibilities. Again, one has a choice about what one focuses one's mind on, when exactly to start the final ascent, what shoes to wear, but prudent or appropriate decisions about these matters (and the associated alternative possibilities) do not seem to be the basis of the meaningfulness of the successful ascent.

Of course, the dialectical situation, abstractly considered, is similar to the situation with respect to the Frankfurt cases. Notoriously, it is very difficult to construct a Frankfurt-style case in which it is both true that the relevant agent is intuitively deemed morally responsible for his behavior and he also has *no* alternative possibilities. Mercifully, one does not have to sort through the voluminous and often complex literature on Frankfurt-style cases to see the point: even if there are residual alternative possibilities in the Frankfurt-style cases, at least in some versions these alternative possibilities are *irrelevant* to the agent's moral responsibility. Like the possibility of a pilgrim to begin her journey at 7:05 a. m. rather than 7:00 a. m. these residual alternative possibilities are what I have called "mere flickers of freedom". They are insufficiently robust to make a difference to the agent's moral responsibility. Mere flickers of freedom are shadowy—even ghostly—things. If (and that is a big "if") you are inclined to hold that moral responsibility requires genuine access to alternative possibilities, these are not the sort of alternative possibilities you have (or ought to have) in mind. So next time you read an article with the "nth" iteration of a Frankfurt-style case in which the author alleges that he has shown that there are indeed residual alternative possibilities, ask yourself: Does it really matter? Are these the sorts of alternative possibilities that can plausibly ground moral responsibility, on the picture according to which we do indeed require alternative possibilities for moral responsibility? Or are they mere flickers of freedom?

Thus, the Frankfurt-style cases, much reviled by some, point to something deep and important: moral responsibility depends on the features (perhaps modal or dispositional) of the actual path—the actual sequence—and *not* on the availability of alternative paths. The cases strongly suggest, but admittedly do not (even with ancillary argumentation) *entail* this result. But it is unreasonable to suppose that *any* philosophical thought-experiment will uncontroversially *entail* a substantial and contentious result in a great philosophical debate. Again: the next time you read an article in which the author triumphantly declares that he has found a reason to think that the Frankfurt-style cases, no matter how much you fancy them up, fall slightly short of entailing the intended conclusion, ask yourself: Why should anyone be surprised by this at all? Isn't the point of the cases to latch onto and resonate with some basic intuitive judgments, thus strongly suggesting a new way of conceptualizing moral responsibility—a new paradigm? Is it *ever* the case in philosophy that a new paradigm is actually *entailed* by the evidence for it?

VI. Conclusion

Over the years (and I hope I still have many to go) I have attempted to fill in the actual-sequence model of moral responsibility. Even though it does not matter for moral responsibility whether we have freedom to do otherwise, it *does* matter whether we act freely or exhibit a distinctive kind of control, which I have called "guidance control". I, sometimes together with my co-author and friend, Mark Ravizza, have given an account of guidance control on which such control is compatible with both causal determinism and causal indeterminism. Guidance control, we have argued, is the freedom required for moral responsibility.

The value of exhibiting this distinctive kind of control is not the value of making a difference (as can be seen in the Frankfurt-style cases). Rather, the value is the value of making a statement—of expressing oneself in a certain way. More specifically, when one exhibits guidance control, one writes a part of the narrative of one's life. It is our freedom—our capacity for guidance control—that renders us authors of the narratives of our lives. It is this freedom that transforms our lives into the subject matter for genuine narratives and that makes it the case that our lives have an irreducibly narrative dimension of value.