JOHN MARTIN FISCHER

RESPONSIBILITY, CONTROL, AND OMISSIONS

(Received and accepted 28 September 1995)

ABSTRACT. Previously, I have argued that moral responsibility for actions is associated with guidance control. This sort of control does not necessarily involve the freedom to do otherwise. In this paper I extend the view to apply to omissions. That is, moral responsibility for an omission is associated with guidance control of that omission. This helps to provide a systematic, unified account of moral responsibility.

KEY WORDS: control, free will, moral responsibility, omissions

I. OMISSIONS AND ALTERNATIVE POSSIBILITIES

I believe that an individual can be held morally responsible for an action, even though he could not have done otherwise (Also, I believe that there are examples in which one can be held morally responsible for a consequence, in which the agent could not have prevented the consequence from obtaining). Actions and consequences are instances of what might be called "positive agency." It will be useful to have before us here an example involving positive agency.

Matthew is walking along a beach, looking at the water. He sees a child struggling in the water, and he quickly deliberates about the mattter, jumps into the water, and rescues the child. We can imagine that Matthew does not give any thought to not trying to rescue the child, but that if he had considered not trying to save the child, he would have been overwhelmed by literally irresistible guilt feelings which would have caused him to jump into the water and save the child anyway. I simply stipulate that in the alternative sequence the urge to save the child would be genuinely irresistible.

Apparently, Matthew is morally responsible – indeed, praiseworthy – for his action, although he could not have done otherwise. Matthew acts freely in saving the child; he acts exactly as he would have acted, if he had lacked the propensity toward strong feelings of guilt. Here is a case in which no responsibility-undermining factor operates in the actual sequence and thus Matthew is morally responsible for what he does (A "responsibility-undermining factor" rules out moral responsibility; intuitively, one tends to think of certain kinds of brainwashing, hypnosis, subliminal advertising,

The Journal of Ethics **1**: 45–64, 1997.

coercion, and direct manipulation of the brain as such factors). And yet the presence of Matthew's propensity toward very strong feelings of guilt renders it true that Matthew could not have done otherwise. Call the case of Matthew, "Hero"[1] (Of course, one could alter the case so that there would be a "counterfactual intervener" – such as a nefarious or even nifty and nice neurologist – associated with Matthew, as in the typical Frankfurt-type examples).

Hero is just one example of many cases of moral responsibility for positive agency in which the agent does not have alternative possibilities. Are there such cases – in which the agent does not have alternative possibilities – of moral responsibility for negative agency (i.e., for omissions)?

Before considering some examples, let me pause to say a few words about the admittedly very problematic notion of "omissions." There are various different conceptions of omissions. One way of classifying them distinguishes wider and narrower conceptions of omissions. On the wider conception (which may not link up closely with ordinary usage), whenever a person does not do something, *X*, he fails in the relevant sense to do it, and he omits to do it. Thus, we are all now failing to stop the Earth's rotation (and omitting to stop the Earth's rotation). Omission to do *X* (according to the wide conception) need not require explicit deliberation about *X*, and it need not require the ability to do *X*. I shall, in part for the sake of simplicity, adopt this wide conception of omissions. My views, however, are compatible with various ways of narrowing the notion of omissions. And even if one takes a rather narrow view of what an omission is, it still is important to have an account of moral responsibilty for failures which don't count as omissions (narrowly construed); after all, in ordinary usage we do talk of moral responsibilty for not doing *X* (where this not-doing may not count as an omission, narrowly construed).

Consider, now, an example I shall call "Sloth."[2] In "Sloth," 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.

Is John morally responsible for failing to save the child? Unbeknownst to John, the child was about to drown when John glimpsed him, and the child drowned one second after John decided not to jump into the water. I believe that the facts of the case exert pressure to say that John is not morally responsible for failing to save the child: after all, the child would

[1] The case is presented and discussed in J.M. Fischer and M. Ravizza, "Responsibility and Inevitability," *Ethics* 101 (1991), pp. 258–278.

[2] The examples presented in this section are from Fischer and Ravizza, 1991.

have drowned, even if John had tried to save it. John could not have saved the child. John may well be morally responsible for deciding not to try to save the child and even for not trying to save the child, but he is *not* morally responsible for not saving the child. "Sloth" is no different in this respect from a case ("Shark") exactly like it except that the child would not have drowned immediately; rather, a patrol of sharks which (unbeknownst to John) infested the water between the beach and the struggling child would have eaten John, had he jumped in.

In Sloth and Shark, it seems clear that John cannot fairly be held morally responsible for failing to save the child. This of course is compatible with John's being appropriately held responsible for deciding not to try to save the child, for not trying to save the child, and so forth. It's just that he cannot appropriately be held morally responsible for *not saving the child.*

Imagine, similarly, that Sue thinks that she can end a terrible drought by doing a rain dance. Of course, we would say that Sue (although quite sincere in her convictions) does not in fact have the power to affect the weather. Suppose, also, that there are no clouds in sight (and no clouds within hundreds of miles); atmospheric conditions imply that it will not rain for weeks. Now Sue happens to hate the local farmers, and she would like to hurt them in any way possible. While falsely believing that she could easily end the drought immediately, she deliberately refrains from doing her rain dance.[3] Let us call this example, "Rain Dance."

Is Sue morally responsible for failing to cause it to rain (i.e., for not ending the drought) in "Rain Dance"? Again, there is pressure to say that, whereas Sue might be morally responsible for not doing the rain dance and for not trying to end the drought, she is *not* morally responsible for not ending the drought. After all, Sue could not have ended the drought.

The cases presented above are cases in which an agent omits to do something *good.* I now turn to a similar case in which an agent omits to do something *bad.* Imagine that you are a small-time thug strolling along a dimly-lit street in a deserted part of town. Suddenly you spy a shiny, new Mercedes with a flat tire stranded by the side of the road. The driver of the car is a well-dressed, elderly gentleman with a bulging billfold in his breast pocket. You are tempted to hurry over to the car, assault the old man, and steal his money. Fortunately, you decide against this, and you continue along your way.

Are you morally responsible for failing to rob the driver? Well, unbeknownst to you (and the driver of the car), the Mafia has put drugs into the trunk of the car. Five Mafioso thugs are watching the car from five other cars in the neighborhood. They have strict instructions: if anyone threatens

[3] This kind of example is due to C. Ginet.

the driver of the car, they are to shoot him with their "Uzi's." In these circumstances, we can safely imagine that, if you had attempted to rob the driver, you would have been killed.

I believe that you are *not* morally responsible for failing to rob the driver. You might be morally responsible for deciding not to rob the driver, for not deciding to rob the driver, and for not trying to rob the driver. But there is strong pressure to say that you are simply *not* morally responsible for not robbing the driver, and this pressure seems to come from the fact that you *could not* rob the driver.

These cases suggest that an agent cannot be held morally responsible for not performing an action which he cannot perform. Thus, these cases in conjunction with "Hero" (and a whole array of cases of positive agency) suggest that actions and omissions are asymmetric with respect to the requirement of alternative possibilities. That is, it seems that moral responsibility for an action does not require the freedom to refrain from performing the action, whereas moral responsibility for failure to perform an action requires the freedom to perform the action. A similar asymmetry is suggested for moral responsibility for consequences and moral responsibility for omissions: moral responsibility for a consequence does not require the freedom to prevent the consequence from occurring, whereas moral responsibility for failure to perform an action requires the freedom to perform the action. Although the "asymmetry thesis" holds that positive agency in general – actions and their consequences – is relevantly different from omissions with respect to the requirement of alternative possibilities, it will be simpler, especially at first, to focus on the asymmetry between actions and omissions.

II. OMISSIONS AND FRANKFURT-TYPE CASES

The cases of omissions presented in the previous section suggest that moral responsibility for the failure to do *X* requires the ability to do *X*. But there are other cases which suggest precisely the opposite. Here are some.

Consider the following remarks by Harry Frankfurt:

> In "Sloth" ["Shark"] John decides against saving a drowning child who (because there are sharks nearby) would have drowned even if John had tried to save him. Fischer and Ravizza suggest that it is discordant to insist that in these circumstances John is morally responsible for not saving the child. They are right about this. But what explains the discordance is not, as they suppose, the fact that it was impossible for John to save the child.
>
> This fact might have been due to circumstances of quite a different sort than those that they describe. Thus, imagine that if John had even started to consider saving the child, he would have been overwhelmed by a literally irresistible desire to do something else; and imagine that this would have caused him to discard all thought of saving the child. With

this change, the case of John exactly parallels another of Fischer's and Ravizza's examples – that of Matthew ("Hero").[4]

In virtue of the apparent parallel status of Hero and the Frankfurt-style version of Sloth, Frankfurt holds that John should be considered morally responsible for failing to save the child (in his version of Sloth). If so, this is a case in which an individual is morally responsible for failing to do *X* even though he *cannot* do *X*.

Other philosophers have presented similar "Frankfurt-type" omissions cases.[5] Clearly, the Frankfurt-type version of Sloth could be developed with a counterfactual intervener. Here is just this sort of case (developed by Randolph Clarke):

> Sam promises to babysit little Freddy. But Sam forgets. No one makes Sam forget; it just slips his mind. Consequently, he fails to show up to babysit little Freddy. Unbeknownst to Sam, a mad scientist is monitoring his thoughts. Had Sam been going to remember his promise, the scientist would have intervened and prevented him from remembering it. The scientist would not have intervened in any other way. As it happened, the scientist did not intervene at all; there was no need to.[6]

Clarke's analysis of this case, call it "Babysitter," is as follows:

> Here ... Sam's not showing up depends on his forgetting; had Sam remembered, nothing would have prevented him from keeping his promise. He would have done so. And Sam is responsible for forgetting. Since his not showing up *depends* in this way on something for which he is responsible, it seems to me that he is responsible for not showing up.[7]

Clarke goes on to suggest a principle according to which an agent is morally responsible for an omission to perform a certain action only if: had she intended to perform that action, and had she tried to carry out that intention, then she would have performed the omitted action.

A very similar view about moral responsibility for omissions is defended by Alison McIntyre.[8] She first presents the following case, which appears to confirm the idea (of the previous section) that moral responsibility for omissions requires the ability to do the relevant action:

---

[4] H. Frankfurt, "An Alleged Asymmetry between Actions and Omissions," *Ethics* 104 (1994), pp. 620–623, esp. p. 620.

[5] For interesting and useful discussions of moral responsibility for omissions, including Frankfurt-type omissions cases, see: I. Haji, "A Riddle Regarding Omissions," *Canadian Journal of Philosophy* 22 (1992), pp. 485–502; R. Clarke, "Ability and Responsibility for Omissions," *Philosophical Studies* 73 (1994), pp. 195–208; "D. Zimmerman, "Acts, Omissions and 'Semi-compatibilism'," *Philosophical Studies* 73 (1994), pp. 209–223; A. McIntyre, "Compatibilists Could Have Done Otherwise: Responsibility and Negative Agency," *Philosophical Review* 103 (1994), pp. 453–488; and W. Glannon, "Symmetrical Responsibility," *Journal of Philosophy* 92 (1995), pp. 261–274.

[6] Clarke, p. 203.

[7] Clarke, pp. 203–204.

[8] McIntyre, 1994.

You are a forest ranger and a large forest fire is approaching from the north. You believe that you could start a backfire heading north which would burn the timber in the fire's path and thereby prevent the forest fire from continuing southward. More specifically, you believe that you could use the gasoline in your truck's fuel tank and some dry matches in your kitchen to do this. But you decide not to start a backfire, the forest fire sweeps onward, and a large area of forest to the south is destroyed. Unbeknownst to you, the truck's fuel tank has sprung a leak and is now empty, and your matches are sitting in a puddle of water. You couldn't have started a backfire if you had tried. If we suppose that there was no other method of stopping the fire available to you, it follows that you could not have prevented the fire from continuing southward if you had tried.[9]

McIntyre goes on to give this version of the case:

Case 1. It is your duty as a forest ranger to start a backfire and you believe that you should do so, but out of laziness rationalized with the vain hope that the fire will burn itself out, you do nothing to stop the fire. When you come to be aware of what you believe to be the full consequences of your omission you feel terrible.[10]

Here, in "Forest Ranger 1," it seems that you are not morally responsible for failing to start a backfire or for failing to stop the forest fire. McIntyre agrees with this view, but she now presents a Frankfurt-type version of her case; in this version of the case, she assumes that the fuel tank has *not* sprung a leak and the matches are *not* wet, and she says:

You, the forest ranger, decide not to start a backfire to prevent the forest fire from advancing southward. A group of fanatical environmentalists who are zealous opponents of forest fire prevention efforts have hired a super-skilled neurologist to monitor your deliberations. If you had shown any sign of seriously considering the option of starting a backfire, the neurologist would have intervened and caused you to decide not to take any preventive action. As things turned out, you decided 'under your own steam' not to act, but because of the neurologist's monitoring, you could not have decided to start a backfire if you had believed that there was reason to do so, and because of this fact, you could not have started a backfire.[11]

As McIntyre points out, in contrast to her first case, in this case (which I shall call "Forest Ranger 2"), you seem to be morally responsible for failing to start the backfire and thus for failing to stop the forest fire. And this is so, even though you could not have started a backfire and you could not have stopped the forest fire. You are responsible for your failures here, on her view, because in the Frankfurt-type version of the case, "you could have started a backfire [and thus stopped the forest fire] if you had decided to do so and had tried."[12] This fact highlights the difference between Forest

[9] McIntyre, p. 458.
[10] McIntyre, p. 458.
[11] McIntyre, pp. 465–466.
[12] McIntyre, p. 466.

Ranger 1 and 2: in Forest Ranger 1 you would not have started a backfire if you had tried (because of the leaking fuel tank and wet matches); but in Forest Ranger 2, you would have suceeded in starting a backfire, if you had tried.

McIntyre and Clarke thus hold a similar view: they contend that in cases in which one could have performed the relevant action, if one had decided (and/or tried), one can be morally responsible for the omission. That is, McIntyre and Clarke hold that when one's ability to do the act in question is dependent upon one's decision (and/or efforts), then one may be morally responsible for failing to do *X*, even if one cannot do *X*.

Reflection on the cases of omissions presented in this and the previous section leads to a puzzle. Cases such as Sloth, Shark, Rain Dance, and Flat Tire render it plausible that in order to be morally responsible for failing to do *X*, one must be able to do *X*. However, cases such as the Frankfurt-style Sloth case, Babysitter, and Forest Ranger 2 suggest precisely the opposite. If one wants to say what seems plausible about the Frankfurt-style omissions cases, how can one also say what is plausible about the first range of cases?

There are cases of positive agency – performing actions and bringing about consequences of those actions – in which moral responsibility does not require alternative possibilities. But in the realm of negative agency – omissions – we have a puzzle: in part of the realm it seems that there is a requirement of alternative possibilities for moral responsibility, but in another part of the realm it seems that there is no such requirement.

I believe the puzzle can be solved by appeal to an association of moral responsibility with control. In the following section, I shall present (in an admittedly very sketchy fashion) some tools that will be helpful in seeking to solve the puzzle. Then I shall employ these tools to argue that the conditions for moral responsibility for positive and negative agency are *symmetric*; in neither case does moral responsibility require alternative possibilities. I shall maintain that there *is indeed* an interesting difference between the two groups of omissions cases described above; but I shall show how this difference can be acknowledged compatibly with the view that moral responsibility for neither positive nor negative agency requires alternative possibilities.

III. SOME TOOLS TO SOLVE THE PUZZLE

1. *Two Kinds of Control*

It seems to me that the conclusion tentatively adopted above about positive agency is correct: moral responsibility for positive agency does not

require the sort of control which involves the existence of genuinely open alternative possibilities. But this is not to say that moral responsibility in the context of positive agency does not require control of *any* sort. Indeed, it is important to distinguish two sorts of control, and it will emerge that moral responsibility for positive agency is associated with one (but not the other) kind of control.[13]

Let us suppose that I am driving my car.[14] It is functioning well, and I wish to make a right turn. As a result of my intention to turn right, I signal, turn the steering wheel, and carefully guide the car to the right. Further, I here assume that I was able to form the intention *not* to turn the car to the right but to turn the car to the left instead. Also, I assume that had I formed such an intention, I would have turned the steering wheel to the left and the car would have gone to the left. In this ordinary case, I guide the car to the right, but I could have guided it to the left. I control the car, and also I have a certain sort of control *over* the car's movements. Insofar as I actually guide the car in a certain way, I shall say that I have "guidance control." Further, insofar as I have the power to guide the car in a different way, I shall say that I have "regulative control" (Of course, here I am not making any special assumptions, such as that causal determinism obtains or God exists).

To develop these notions of control (and their relationship), imagine a second case. In this analogue of the Frankfurt-type case presented above, I again guide my car in the normal way to the right. The car's steering apparatus *works properly* when I steer the car to the right. But unbeknownst to me, the car's steering apparatus is broken in such a way that, if I were to try to turn it in some other direction, the car would veer off to the right in precisely the way it actually goes to the right.[15] Since I actually do not try to do anything but turn to the right, the apparatus functions normally and the car's movements are precisely as they would have been, if there had been no problem with the steering apparatus. Indeed, my guidance of the car to the right is precisely the same in this case and the first car case.

Here, as in the first car case, it appears that I control the movement of the car in the sense of guiding it (in a certain way) to the right. Thus,

---

[13] For a parallel distinction between two kinds of control, see M.J. Zimmerman, *An Essay on Moral Responsibility* (Totowa, NJ: Rowman and Littlefield, 1988), pp. 32–34.

[14] For this following discussion, see J.M. Fischer, *The Metaphysics of Free Will: An Essay on Control* (Cambridge, MA: Blackwell, 1994), pp. 132–134.

[15] Note that the example would have precisely the same implications if alternative possibilities were ruled out by virtue of the existence of *another agent*. So imagine that the car is a "driver instruction" automobile with dual controls. Although I actually guide the car to the right, we can imagine that the instructor could have intervened and caused the car to go to the right, if I had shown any inclination to cause it to go in some other direction.

I have guidance control of the car. But I cannot cause it to go anywhere other than where it actually goes. Thus, I lack regulative control of the car. I control the car, but I do not have control *over* the car (or the car's movements). Generally, we assume that guidance control and regulative control go together. But this Frankfurt-type case shows how they can at least in principle pull apart: one can have guidance control without regulative control. That is, one can have a certain sort of control without having the sort of control that involves alternative possibilities (The Frankfurt-type cases of actions and omisisons presented above have this structure. For example, in Hero Matthew has guidance control of his saving the child, even though he lacks regulative control over his saving the child. And so on).

The Frankfurt-type cases (involving actions), unusual as they are, may well point us to something as significant as it is mundane. When we are morally responsible for our actions, we *do* possess a kind of control. So the traditional assumption of the association of moral responsibility (and personhood) with control is quite correct. But it need not be the sort of control that involves alternative possibilities. The suggestion, derived from the Frankfurt-type cases, is that the sort of control necessarily associated with moral responsibility for action is *guidance control.* Whereas we may intuitively suppose that regulative control always comes with guidance control, it is not, at a deep level, regulative control that grounds moral responsibility.

I have not sought to give a precise (or even very informative) account of the two sorts of control. Rather, I have relied on the intuitive idea that there is a sense of control in which I control the car when I guide it (in the normal way) to the right. Further, I have employed the Frankfurt-type example to argue that this sense of control need not involve any alternative possibilities. Then, I have simply contrasted this sort of control with a kind of control which does indeed require alternative possibilities. Now I shall attempt to say more (just a bit more!) about the first sort of control – guidance control. It is this sort of control which, I have claimed, is associated with moral responsibility for actions.

2. *Guidance Control of Actions*

The basic idea is that an agent has guidance control of an action insofar as the action issues from the agent's own moderately reasons-responsive mechanism. Although what it is for a mechanism to be an "agent's own" is a difficult and important issue, I shall not be addressing it here; here

I shall simply assume an intuitive understanding of this notion.[16] Further, I shall here give a only the scimpiest sketch of moderate reasons-responsiveness.[17] To say whether an action issues from a moderately reasons-responsive mechanism, we first need to identify the kind of mechanism that actually issues in action. It is important to see that, in some cases, intuitively different kinds of mechanisms operate in the actual sequence and the alternative sequence. So, for instance, in "Hero," Matthew's actual-sequence mechanism is of a different sort from his alternative-sequence mechanism; in the actual sequence, he quickly deliberates and decides to save the struggling child, and his reasoning is uninfluenced by any overwhelming urge. However, in the alternative sequence, his deliberations are influenced by an overwhelming and irresistible urge to save the swimmer. Whereas it is difficult to produce an explicit criterion of mechanism-individuation, I believe that it is natural to say that in Frankfurt-type cases different sorts of mechanisms issue in the actions in the actual and alternative sequences; indeed, this seems to be definitive of Frankfurt-type cases.

For a mechanism to be moderately responsive to reasons, it must at least be weakly reasons-responsive. In order to determine whether a mechanism of a certain type is weakly reasons-responsive, one asks whether there exists some possible scenario (with the same natural laws as the actual world) in which that type of mechanism operates, the agent has reason to do otherwise, and the agent does otherwise (for that reason). That is, we hold fixed the actual type of mechanism (and natural laws), and we ask whether the agent would respond to *some* possible incentive to do otherwise. Now moderate responsiveness differs from mere weak responsivness in that it demands not only that the agent would respond to at least one possible incentive to do otherwise; the agent must also exhibit a minimally coherent pattern of *recognition* of reasons. If (under the envisaged circumstances) the agent would so respond, then the actually operative mechanism is moderately reasons-responsive. In contrast, strong reasons-responsiveness obtains when a certain kind of mechanism (K) actually issues in an action and if there were sufficient reason to do otherwise and K were to operate, the agent would recognize the sufficient reason to do otherwise and thus choose to do otherwise and do otherwise.

Let me say a little more about strong, weak, and moderate reasons-responsiveness. Under the requirement of strong reasons-responsiveness,

---

[16] For an attempt to say what it is for a mechanism to be "the agent's own," see J.M. Fischer and M. Ravizza, *Responsibility and Control: A Theory of Moral Responsibility* (manuscript; forthcoming, Cambridge: Cambridge University Press).

[17] For further details, see Fischer and Ravizza, forthcoming.

we ask what would happen if there were a sufficient reason to do otherwise (holding fixed the actual kind of mechanism). Strong reasons-responsiveness points us to the alternative scenario in which the actual kind of mechanism operates and there is a sufficient reason to do otherwise which is most simliar to the actual situation. Strong reasons-responsiveness is similar to Robert Nozick's notion of "tracking value" or "tracking bestness."[18] In contrast, under weak reasons-responsiveness, there must simply exist *some* possible scenario in which the agent's actual kind of mechanism operates, the natural laws are held fixed, there is a sufficient reason to do otherwise, and the agent does otherwise. A weak-willed agent may exhibit weak-reasons responsiveness, even though he does not exhibit strong-reasons responsiveness (similarly for a morally bad agent). Finally, moderate reasons-responsiveness is stronger than weak reasons-responsiveness, and weaker than strong reasons-responsiveness. Whereas it demands only the sort of reactivity to reasons posited by weak reasons-responsiveness, it demands more in the way of reasons-recognition. It thus demands a certain sort of "normative competence."[19]

3. *Guidance Control of Consequences*

The account of guidance control of consequences is in certain respects parallel to (and also an extension of) the account of guidance control of actions. The leading idea is that the agent displays guidance control of a consequence insofar as the consequence emanates from a responsive *sequence*. It is necessary, in the context of a consequence that is more than simply a bodily movement, to distinguish *two components* of the sequence leading to the consequence. The first component is the mechanism leading to the bodily movement, and the second component is the process leading from the bodily movement to the event in the external world. I shall say that, in order for the sequence leading to a consequence to be responsive, both the mechanism leading to the bodily movement must be moderately reasons-responsive and the process leading from the bodily movement to the event in the external world must be "sensitive to the bodily movement."

Before proceeding, it is important to note that the counterfactual intervener in a Frankfurt-type case need not be another agent (whose action in the alternative sequence would bring about the consequence in question). As Frankfurt points out, the role of counterfactual intervener may be played

[18] R. Nozick, *Philosophical Explanations* (Cambridge, MA: Harvard University Press, 1981).

[19] For a development of the relevant notion of normative competence, see Fischer and Ravizza, *Responsibility and Control: A Theory of Moral Responsibility*, forthcoming.

"by natural forces involving no will or design at all."[20] It seems, then, that in evaluating the sensitivity of a process one wants to hold fixed not only the actions of other agents in the actual sequence, but also any natural events which play no role in the actual sequence but which would, in the alternative sequence, *trigger* causal chains leading to the consequence in question. For convenience we can group *both* other actions that would trigger causal chains leading to the consequence *and* natural events that would do so under the heading, "triggering events." Let us think of a triggering event (relative to some consequence *C*) as an event which is such that if it were to occur, it would *initiate* a causal sequence leading to *C*.

Now I can present the account of guidance control of consequences as follows. As I have said above, the bodily movement must be moderately reasons-responsive. Further, the process leading from the bodily movement to the event in the external world must be "sensitive to bodily movement" in roughly the following sense: if the actual type of process were to occur and all triggering events which do not actually occur were not to occur, then a different bodily movement would result in a different upshot. Guidance control of a consequence then involves two interlocked – and linked – sensitivities.

Here is a bit more explicit statement of the account. Suppose that in the actual world an agent *S* moves his body in way *B* via a type of mechanism *M*, and *S*'s moving his body in way *B* causes some consequence *C* via a type of process *P*.[21] I shall say that the sequence leading to the consequence *C* is responsive if and only if there exists some way of moving *S's* body *B** (other than *B*) such that: (i) there exists some possible scenario in which an *M*-type mechanism operates, the agent has reason to move his body in way *B**, and the agent does move his body in *B**; and (ii) if *S* were to move his body in way *B**, all triggering events which do not actually occur were not to occur, and a *P*-type process were to occur, then *C* would not occur.

---

[20] Frankfurt, 1969, note 4.

[21] This is still only an approximation to an adequate account. For example, it assumes that there is an appropriate range of scenarios in which *S* recognizes reasons to move his body in way *B**. It also assumes that in the alternative possible scenario the agent moves his body in way *B** for the relevant reason *qua* reason, and so forth.

I shall here also assume that there is just one causal sequence leading to the consequence; thus, in this paper I am concerned with cases of "pre-emptive overdetermination" rather than "simultaneous overdetermination." Further, the focus here is on what might be called "action-triggered" consequences. There might also be "omission-triggered" consequences for which an agent might be morally responsible. A more complete theory of responsibility – one which attends to the full range of possible cases – is presented in Fischer and Ravizza *Responsibility and Control: A Theory of Moral Responsibility*, forthcoming.

Let me now take a moment to discuss a few points which should help both to clarify and to illustrate the principle. (1) In formulating the definition of a responsive sequence, I make use of the intuitive notion of a "type of process" leading from the bodily movement to the event in the external world. This is parallel to the notion of a kind of mechanism issuing in action. I concede both that process-individuation is problematic and that I do not have an explicit theory of process-individuation. But I believe that there is a relatively clear intuitive distinction between different types of processes, just as there is a relatively clear intuitive distinction between different kinds of mechanisms leading to bodily movements.

I do not deny that there will be difficult questions about process-individuation. Nevertheless, all that is required for my purposes here is that there be agreement about some fairly clear cases. If we are unsure about an agent's moral responsibility for a consequence in precisely those cases in which we are unsure about process-individuation, then at least the vagueness in our theory will match the vagueness of the phenomena it purports to analyze.

(2) In ascertaining the responsiveness of a particular sequence involving a mechanism issuing in a bodily movement, a bodily movement, and a process leading from that bodily movement to a consequence, we "hold fixed" the actual type of mechanism and the actual type of process. If it is the case that a different mechanism or process would have taken place if things had been different (i.e., if the case is a Frankfurt-type case), this is irrelevant to the responsiveness of the *actual* sequence.

Further, imagine that we are testing the sensitivity of a particular process leading from a bodily movement to a consequence. Suppose that the agent actually moves his body in a certain way thus causing some consequence, and that no one else actually performs that type of action. Under these conditions, we "hold fixed" others' behavior when we test for the sensitivity of the process leading from action to consequence. The point is that, when we are interested in the sensitivity of the process to action, we are interested in whether there would have been a different outcome, if the agent had not performed a certain sort of action *and all non-occurring triggering events were not to occur.*

The sequence leading to a consequence (of a certain sort) includes more than just the mechanism issuing in bodily movement. Thus, both components – i.e., the mechanism leading to the bodily movement and the process leading from the bodily movement to the event in the external world – are relevant to guidance control of a consequence. The account of guidance control of a consequence involves what might be called "two stages."

(3) The notion of a "triggering event" is – like the notions of "mechanism" and "process" – fuzzy around the edges. But, again, I believe that it is tolerably clear for the present purposes. Note that a triggering event is an event which would "inititate" a causal chain leading to a certain consequence. Although the concept of "initiation" is difficult to articulate crisply, we rely on the fact that there are some fairly uncontroversial instances of the concept. So, for example, if a lightning bolt hits a house and there is a resulting fire, the event of the lightning's hitting the house could be said to initiate the sequence leading to the destruction of the house. And this is so even if there were certain atmospheric events which antedated the lightning bolt and which led to it. Of course, the notion of "initiation" is highly context-dependent, and the truth of claims about purported initiations will depend on the purposes and goals of the individuals making (and considering) the claims. But I believe that the notion of initiation issues in tolerably clear intuitive judgments about the cases relevant to our purposes.

IV. THE SYMMETRIC PRINCIPLE OF MORAL RESPONSIBILITY

Now the tools for resolving the puzzle about moral responsibility for omissions are at hand. I have suggested that in cases of positive agency, moral responsibility is associated with control in a certain way. More specifically, I have claimed that guidance control is the kind of control associated with moral responsibility in cases of positive agency. I started with actions, and developed an account of guidance control of actions. This account employs the notion of moderate reasons-responsiveness. I then built on this model to develop an account of guidance control of consequences. On this account, there may be two steps: the bodily movement and then some event in the external world. In order for the sequence (involving both steps) to be appropriately responsive to reason, the bodily movement must be moderately reasons-responsive, and the event in the external world must be sensitive to the bodily movement.

The key to resolving the puzzle about omissions is to develop an analogous account of guidance control for omissions. If guidance control is all the control required for moral responsibility for omissions, then perhaps one can say just the right thing about the entire array of cases presented in the first two sections of this paper. On this approach, it is *not* the case that alternative possibilities are required for any part of the realm of omissions. Whereas it may seem that the only way to explain why an agent is not morally responsible for certain omissions is to cite his inability to perform the relevant action, another explanation is available: the agent may lack

guidance control of the omission. Further, on this approach positive and negative agency are symmetric with respect to the requirement of alternative possibilities: guidance control (and not regulative control) is the kind of control associated with moral responsibility for positive and negative agency. Let us call this the "Symmetric Principle of Moral Responsibility" (in the rest of this paper, I shall be focusing primarily on the negative-agency component of the Symmetric Principle; thus, when I speak of the Symmetric Principle, I shall be speaking about the component of it which claims that guidance control is the sort of control necessary and sufficient for moral responsibility for omissions).

Like actions (and their consequences), omissions may be relatively simple or complex. A simple omission would be the failure to move one's body in a certain way (Let us call these "bodily omissions"). In these cases, the failure to move one's body in a certain way "fully constitutes" the omission.[22] Here the application of the notion of guidance control is also relatively simple: it is natural to say that one has guidance control of one's failure to do *X* (in a case of a bodily omission) just in case one's failure to do *X* issues from one's own, moderately reasons-responsive mechanism. As with the case of actions, one here holds fixed the actual-sequence mechanism that issues in the failure to move one's body in a certain way, and asks what would happen in a relevant range of alternative scenarios. The account is parallel to the account in the case of action.

A bit more specifically, let us suppose that the failure to do *X* here is the failure to move one's body in a certain way *B** which actually occurs via mechanism *M*. What is it for one's failure to move one's body in way *B** to issue from a moderately reasons responsive mechanism? It must be the case that, if *M* were to operate and the natural laws were held fixed, there is at least some scenario in which one has reason to move in way *B** and one does so (for that reason).

The treatment of moral responsibility for more complex omissions – omissions that are not simply bodily omissions – is analogous to the treatment of moral responsibility for the consequences of one's actions. As I pointed out above, in the context of assessing moral responsibility for the consequences of one's actions, there are typically two steps or stages: the bodily movement must be moderately responsive to reasons, and the event in the "external world" must be appropriately sensitive to one's action. Just as one holds fixed the actual-sequence mechanism when assessing the

[22] These omissions are like Frankfurt's "personal" failures; see H. Frankfurt, "What We Are Morally Responsible For," in L.S. Cauman, I. Levi, C. Parsons, and R. Schwartz (eds.), *How Many Questions? Essays in Honor of Sidney Morgenbesser* (Indianapolis, IN: Hackett, 1982), pp. 321–335.

moderate reasons-responsiveness of the bodily movement, one holds fixed the actual conditions in the world when assessing the sensitivity of the external event to the bodily movement. That is, at both stages one holds fixed the relevant features of the actual sequence; and at both stages one looks for a certain sort of responsiveness or sensitivity.

In the context of complex omissions, the account of guidance control is parallel to the account of guidance control of consequences of actions. It is natural to say that an agent has guidance control of his failure to do *X* (where *X* is not simply a bodily movement) just in case there exists a way of moving his body (different from the way he actually moves it) such that: 1) his failure to move his body in this way issues from his own, moderately reasons-responsive mechanism, and 2) the relevant event in the external world is suitably sensitive to that failure to move his body. The details of the analyses are understood to be parallel to those in the context of positive agency.

More specifically, what is it for one's failure to do *X* here to issue from a sequence in which the agent has guidance control? Let us suppose that the agent fails to move his body in way *B**, and this issues in some result in the world *C* via process *P* (This counts as the agent's not doing *X*). It must be the case that i) there exists a way of moving his body *B** (different from how he actually moves it) such that, if the actual type of mechanism *M* were to occur and the natural laws were held fixed, then there is some scenario in which there is a reason for the agent to move in way *B** and he does so (for that reason); and ii) if he were to move in way *B**, process *P* were to occur, and all non-occurring triggering events were not to occur, then some different result *C** would occur.

It should be noted that the notion of a mechanism leading to a *failure* (a failure to move one's body, or an omission) is not as clear as the notion of a mechanism leading to an action. To make this notion as clear as possible, I shall say that the mechanism leading to an omission is the mechanism leading to *what the agent does instead.* So, for example, in "Sloth," John walks along the beach instead of jumping in to save the child; thus, I shall say that the mechanism leading to his not jumping in to save the child is the mechanism that leads to his action of walking along the beach.[23]

Let us now apply this account to the range of examples presented above. In "Sloth," "Shark," "Rain Dance," and "Flat Tire," the agents all actually move their bodies in certain ways; thus, they all fail to move their bodies in certain other ways. These failures are plausibly taken to be moderately responsive to reasons. In all these cases, however, there is a problem at the second stage: the relevant events in the external world are not suitably

[23] I am indebted to C. Ginet for this suggestion.

sensitive to the agents' bodily movements (or failures to move their bodies in certain ways). So, in "Shark," John's failure to jump into the water and head toward the struggling child is moderately responsive to reason (he is thus morally responsible for his bodily omission). But even if John had moved his body in this alternative way, the child would have drowned – the sharks would have eaten him (so John is not morally responsible for the complex omission). Similarly, although Sue's failure to do the rain dance is moderately responsive to reason, the drought would not have ended (presumably), even if she had done it. Whereas Sue is responsible for the bodily omission, she is not responsible for the complex omission.

The same sort of analysis applies to all the cases in the first group. In all of these cases the agents are not morally responsible for the relevant complex omissions because they lack guidance control of the omissions. And they lack such control in virtue of their failure to meet the conditions that pertain to the second stage: sensitivity of the external event to one's bodily movements.

Now consider the second group of cases: Frankfurt-type omissions (the Frankfurt-type "Sloth" case, "Babysitter," and "Forest Ranger 2"). In all of these cases the agents lack the ability to do the relevant action. But in all of these cases the agents have guidance control of the relevant omissions, and thus are appropriately considered morally responsible for those omissions.

Take, for example, the Frankfurt-type "Sloth" case. Here, in virtue of his propensity toward strong feelings of guilt, John cannot move his body in any way other than the way he actually does, and thus he cannot save the child. But nevertheless his actual bodily movements issue from a moderately reasons-responsive mechanism. After all, the guilt feelings play no role in the actual sequence – they are not a part of the mechanism that actually issues in action. Further, the child would have been saved (presumably), if John had moved his body in certain different ways. Thus, John's failure to move his body in the relevant way is moderately responsive to reason, and the child's not being saved is sensitive to that failure to move his body. So John has guidance control of his failure to save the child, and is morally responsible for it. And the same sort of analysis applies to all the Frankfurt-type omissions cases.

V. AN OBJECTION

The Symmetric Principle, then, seems to imply all the right judgments about the cases assembled above. But I shall now turn to an objection to it. McIntyre objects with an example and a set of ancillary considerations. First the example:

A meeting of the New York Entomological Society features an international array of dishes prepared using insects. [McIntyre here refers to Maialisa Calta, "Bug Seasoning: When Insect Experts Go in Search of Six-Legged Hors d'oeuvres," *Eating Well* 3 (1992), pp. 22–24.] You, a guest, are invited to sample a tempura dish made of fried crickets. You don't find the prospect of eating insects appealing, though you don't find it disgusting either, and you decline the offer. Suppose that in order to have decided to accept the offer, you would have had to look more closely at the fried crickets. But if you had looked more closely you would have been overwhelmed with revulsion and would have been incapable of deciding to eat some. Since you never do look more closely at the crickets, you decide not to have any without experiencing any feelings of revulsion, and without even suspecting that you would feel revulsion if you examined the dish more closely.[24]

McIntyre employs this example, call it "Insects," as part of a critique of the Symmetric Principle. She says:

... this approach, when applied to omissions, would yield too liberal a condition of moral responsibility. It will turn out that you are morally responsible for omitting to eat the crickets even if there is no possible situation in which you, as you actually are disposed and constituted, could have eaten them.[25]

McIntyre's point is that, in the story, you are actually so constituted that you would have been overwhelmed with revulsion if you looked more closely at the fried crickets, and if we assume that this revulsion is so strong that there is no possible situation in which you could have decided to eat the crickets (given this revulsion), it seems implausible to say that you are morally responsible for your failing to eat the crickets.

But recall that as things actually went, the revulsion played absolutely no role in your deliberations and your decision not to eat the crickets. And note that McIntyre's "Insects" case, in the version which she employs to criticize the Symmetric Principle, seems to be precisely parallel to the Frankfurt-style "Sloth" case. Recall that in the Frankfurt-style Sloth case, John fails to save the child and indeed fails to even consider doing so; but if he were to start to consider saving the child, he would have been overwhelmed by a literally irresistible desire to do something else. Here it is Frankfurt's view (and mine) that John is morally responsible for failing to save the child. Since I agree with Frankfurt about his version of "Sloth," and the two cases appear to be parallel, I am inclined to disagree with McIntyre about "Insects." That is, just as John is morally responsible for failing to save the child in Frankfurt's version of "Sloth," so you are morally responsible for failing to eat the crickets in "Insects."

Further support for my position comes from reflection on the theoretical considerations McIntyre invokes as part of her critique of the Symmetric Principle. She says:

[24] McIntyre, pp. 485–486.

[25] McIntyre, pp. 486–487.

... According to that approach [of Fischer and Ravizza], even if *you* could not have decided to eat some crickets because of your propensity to revulsion, *the mechanism* that actually produced your decision could have done so, and, as a result, you can be morally responsible for your omission. Of course, if we can stipulate that you do not have, or are not affected by, your propensity to feel revulsion, then there would be no obstacle to identifying some possible situation in which you eat some crickets. But what justifies this stipulation? It seems that one could quite reasonably object that this is suspiciously similar to inferring that *you could have done otherwise* from the fact that *you could have done otherwise if what would have prevented you from doing otherwise hadn't existed*![26]

I believe McIntyre's criticism here is unfair. On my approach to both actions and omissions, freedom to do otherwise is not required for moral responsibility; rather, what is relevant are features of the actual sequence that leads to the action or the omission. I certainly agree that someone who actually faces some insuperable obstacle to doing otherwise cannot do otherwise, and it would simply be irrelevant, for most purposes, to point out (what might, in any case, be true) that the agent would be able to do otherwise, if the obstacle were subtracted. Since my approach to moral responsibility does not require alternative possibilities, I am not here in the business of assessing an agent's freedom to do otherwise.

Rather, I am interested in evaluating the mechanisms and processess that actually lead to actions, consequences, and omissions. Since in "Insects" the propensity toward revulsion played no role in your decision or bodily movements, it is not part of the mechanism that actually issues in that decision and those bodily movements. Thus, it is irrelevant to the issue of whether that actual-sequence mechanism is responsive to reasons, and thus also to the issue of whether you are morally responsible for your actions. Clearly, it would be inappropriate to subtract the propensity toward revulsion in considering whether you could have done otherwise; but it is not inappropriate to subtract it when considering whether the actual-sequence mechanism that issues in your omission has a certain feature-responsiveness to reasons.

In focusing on the properties of the actual mechanisms and processes that lead to actions, consequences, and omissions, I am seeking to develop what might be dubbed an "actual-sequence" approach to moral responsibility. But notice that the "actual-sequence" properties fixed on by such an approach may indeed be dispositional properties; as such, their proper analysis may involve (for example) other possible worlds. In the context of an actual-sequence approach to moral responsibility, I have argued that it is required that a reasons-responsive mechanism actually operate; then, I have analyzed reasons-responsiveness in terms of other possible worlds. Whereas other possible worlds are relevant to ascertaining whether there is some

[26] McIntyre, p. 486.

actually operative dispositional feature (such as reasons-responsiveness), such worlds are *not* relevant in virtue of bearing on the question of whether some alternative sequence is genuinely accessible to the agent.

VI. CONCLUSION

I started with a puzzle about moral responsibility for failures. Some cases suggest that in order to be morally responsible for failing to do *X*, you must have the ability to do *X*. But other cases suggest exactly the opposite. In response to this puzzle, I have argued that moral responsibility for failing to do *X* does not in fact require the ability to do *X*; rather, it simply requires that the agent have guidance control of his failure to do *X*.

This suggestion has the virtue of treating positive and negative agency *symmetrically*. Further, I can respect the intuitive view that the first group of cases is interestingly different from the second; but the explanation of the difference is not in terms of freedom to do otherwise (or alternative possibilities). The account of guidance control implies that the two groups are different while nevertheless offering a unified account of moral responsibility for omissions. Further, it helps to exhibit a unified, systematic view of moral responsibility for actions, consequences, and omissions. That is, it helps to show how the association of moral responsibility with control – and more specifically, with guidance control – can begin to systematize and illuminate our considered judgments about the *full content* of moral responsibility, which must include responsibility for actions, consequences, and omissions. Finally, the Symmetric Principle helps to establish semicompatibilism; as with positive agency, there is no reason to think guidance control is incompatible with causal determinism, and thus there is no reason to think moral responsibility for omissions is incompatible with causal determinism.[27]

Department of Philosophy
University of California at Riverside
Riverside, CA 92521-0201
USA

[27] This paper has greatly benefited from comments by M. Ravizza; it constitutes the basis for part of the chapter on omissions in our forthcoming book, *Responsibility and Control: A Theory of Moral Responsibility* (Cambridge: Cambridge University Press). Also, I have read this paper, and benefited from comments, at the University of California, Riverside, UCLA, USC, and the University of Arkansas, Fayetteville. My work on this paper has been supported by a National Endowment for the Humanities Fellowship for University Teachers.