

The following material gives more details about the Fischer approach to MR. It comes from my draft textbook, and starts in the middle of a chapter, so it doesn't have a pretty title.

The material is about the theory of guidance control and MR developed by Fischer and Ravizza, and published in their *Responsibility and Control: a Theory of Moral Responsibility*, published 1998 in the Cambridge Studies in Philosophy and Law series, by Cambridge Univ. Press.

This reading goes well with Section 8 of the Fischer chapter in the textbook. In fact, you can ignore Sections 6 and 7 if you find them unhelpful. If you find those Sections useful, then by all means read them carefully. However, we will pay more attention to Section 8 and to the following reading.

-Machina

Theory 3: being responsive to reasons

Out of the numerous recent attempts by philosophers to present an account of responsibility-generating control that people could realistically actually have, there is one that stands out in my mind as being especially promising. As I have been pondering what I personally think about these matters, this account seems to me to have some real advantages. (Do not take my word for it. Look at the arguments and decide for yourself.) This account describes responsibility-generating control as depending on whether a person is responsive in certain specific ways to reasons that pertain to their situation. I will sometimes refer to this account of control as the “reasons-responsive” account. The version presented below comes from the work of John Martin Fischer, a philosopher at the University of California, Riverside, and Mark Ravizza, a former computer engineer who became a professional philosopher and a Jesuit, now on the faculty at Santa Clara University.⁴⁷ This is another compatibilist approach that does not require people have contra-causal free will or ultimate self-determination in order to be morally responsible.

Guidance control. To understand reasons-responsive control, it is probably best to start with something that Fischer and Ravizza refer to as “guidance control.” When you drive a car under normal circumstances, you guide it. You control where it goes. This involves monitoring what is happening, and making adjustments as required. You are aware of what is going on, you understand what needs to be done to change directions, to slow up, to stop, to go again. You make a great many small choices in the process, and carry them out. Sometimes you have to think about what you are doing, but most of the time there is no deliberation. This sort of control might be called “guidance control.” (You guide the car by constant monitoring and adjusting.) Fischer and Ravizza (“F & R”) hold that the responsibility-generating control we have been seeking for some pages is a type of guidance control.

Notice that nothing in the description of guidance control requires contra-causal free will or self-determination. We don’t have to suppose you are uncaused in the choices you make as you drive. In fact, guidance control needs to have those choices caused by what you see, what you know, and the set of abilities you have. You choose to step on the brake when someone suddenly cuts you off. No time for conscious deliberation on the basis of your desires and values. You simply react “instinctively” to the situation. You are in control—guidance control. It doesn’t matter whether your values and beliefs were self-determined or not.

According to F & R, this rough idea of guidance control can be developed to give us a good picture of responsibility-generating control. To do this properly, they suggest we need to focus on the process by which a person comes to perform whatever action they might be

⁴⁷(1998) Another philosopher, Susan Wolf, has proposed a similar view in a highly recognized book, (1990).

thought to be morally responsible for. If that process is of the right sort, where the person exercises his or her own guidance control of the action, then F & R believe we have the basis for the person's moral responsibility for that action.

“Normal” choice-making mechanisms. In order to determine when a person's actions are the result of their own guidance control, rendering them morally responsible for those actions, it is important to distinguish “normal” decision-making processes from processes that ruin a person's guidance control to such an extent that they seem no longer in control of their own actions and thus not morally responsible for them. Hypnotism, “powerful forms of brainwashing and indoctrination, potent drugs, and certain sorts of direct manipulation of the brain” as well as certain very frightening “coercive threats” are examples of ways in which a person's control over their own actions or behavior can be rendered questionable to such an extent that we should say that the person is no longer responsible. As F & R put it, “When persons are manipulated in certain ways, they are like marionettes and are not appropriate candidates for praise or blame.”⁴⁸ All these “responsibility-undermining factors” work to diminish or destroy moral responsibility by diminishing or destroying the person's own normal guidance control mechanisms.

F & R claim that the key difference between the cases in which moral responsibility is undermined and cases in which it is not has to do with the mental mechanisms that are controlling the behavior of the person. The person exercises the responsibility-generating sort of guidance control in “cases in which there is the ‘normal’ unimpaired operation of the human deliberative mechanism”⁴⁹ (I wish they had said “choice-making mechanism,” since we know that *deliberation* often is not involved.) So, when the person is operating “normally,” as far as their mental mechanisms for making choices go, they have the sort of guidance control of the resulting behavior that will support their being morally responsible for it. When the mechanisms are not operating in their “normal, unimpaired” way, then responsibility will be diminished or nonexistent. Enormously frightening coercive threats, potent drugs, and so on are examples of things that can prevent a person's choice-making mechanisms from operating in a “normal, unimpaired” way.

What is meant by “normal, unimpaired” operation of choice-making mechanisms? F & R believe that when people are exercising normal, unimpaired guidance control, their choice-making mechanisms are “responsive” to relevant reasons for acting or refraining from acting. Here is an example based on their discussion: suppose you are successfully hypnotized to make a silly face at the person sitting next to you when the telephone rings. Then when the phone does ring, you will make the face, even though you may have overriding reasons not to do so. (E.g., you want to impress that person with your maturity, and making silly faces will not get the job done.) This illustrates what it is like to *not* be responsive to reasons. A “normal, unimpaired” choice-making process in you *would* be responsive to the overriding reasons for not making the silly face, and so you would in fact not do it.

⁴⁸All quotes in this paragraph from pp. 35 - 36.

⁴⁹Page 36.

So, as a first rough approximation of a theory of responsibility-generating control we have this: a person has the right sort of control over an action to be morally responsible for it when that action is directed by their own reasons-responsive guidance control mechanism.

Reasons-responsiveness. What exactly does it mean for a choice-making process to be responsive to reasons? F & R believe the appropriate definition would be something like this:

A choice-making process counts as being reasons-responsive when the outcome has the *potential* to be affected by reasons the person could recognize and that we can *understand* as coherent and sane (even though they may not be the reasons that would move us)⁵⁰

Suppose you decide to stay up late so you can steal change from your roommate, to buy cigarettes. Let's try to apply the above definition to see if your decision was made by a reasons-responsive choice-making process. To do so, we need to ask, Was your decision-making process so rigid that *even if a sane, coherent reason came along for you not to stay up and steal*, you would still stay up and steal? Unless you are out of your mind, the answer is "no." Surely you would have decided to call it a night if your roommate was seriously dangerous and was suspiciously on guard. Surely you might have made a different choice if you had realized the police were videotaping your every move. There are any number of potential understandable reasons that would have made a difference in your decision. According to the above definition of "reasons-responsive," the existence of these potential scenarios, *in which you would have been responsive to sane, coherent reasons*, shows that your decision-making process is in fact "reasons-responsive."

What is a real-life example of human choice-making that is *not* reasons-responsive? I think one easy example of the sort of thing F & R have in mind comes from obsessive-compulsive disorder. The obsessive-compulsive person feels an inner compulsion to constantly wash hands, or to constantly count things, or to constantly engage in some other patterned ritualistic behavior, no matter what sane and coherent reasons they become aware of to stop. Their choice-making mechanism is simply not responsive to sane, coherent reasons to stop engaging in the compulsive behavior. So, according to F & R's analysis, such people are not responsible for their behavior, since they do not have the right kind of control over it.

We now have a proposal for what kind of control a person has to have in order to be responsible: a person has the right sort of control to be responsible when the process by which they make the relevant choices is reasons-responsive as just defined. (We will have to wait for just a bit before we see how to modify this proposal to account for *moral* responsibility. So far, the proposal is simply a proposal about being in some way responsible.)

⁵⁰From Chapters 2 and 3.

Notice that under this new definition, you are responsible for staying up and stealing change to buy cigarettes, since you presumably would have been dissuaded by any number of sane, coherent reasons. On the other hand, when someone is truly under a strong and irrational psychological compulsion, or under the influence of powerful drugs, *nothing* would be a reason for them to change some particular bit of behavior that results from the compulsion or drug, unless perhaps it would be some bizarre reason that we cannot understand or recognize as sane. This is a sign that they are *out of control*, or, more precisely, that they *lack* reasons-responsive guidance control over that behavior. They will, accordingly, *not* be responsible for that behavior. These are welcome results, for they seem to fit what we would want of a theory about responsibility-generating control.

Why require that the potential reasons to which the person would respond be understandable and sane? If the only reasons that someone would potentially respond to are totally bizarre, disconnected from reality, there is a serious question about whether the person is sane enough about this action to be responsible for it. Suppose you will go into an uncontrollable fit of rage and begin attacking me if I come to visit you in the morning, unless I come precisely at 10:09 am by your watch. There is no rational reason for any of this. 10:09 appears to be a completely arbitrary thing with you. You have nothing against me. 10:09 has no historical significance in your life. So, you are responsive to reasons, in an odd sort of way—my coming at precisely 10:09 would be a reason for you not to attack me. This hardly seems like the sort of reasons-responsiveness that would make you responsible for attacking me when I came at 10:06. Your choice-making process is so messed up that it seems you are not responsible. Instead of saying you deserve punishment for attacking me, we would presumably think you need psychotherapy or sedation. So, we don't want your choice-making process to count as reasons-responsive. What counts as a reason for you is something that we do not see as a genuine reason, since we can't make any sense of it. This illustrates why F & R define reasons-responsiveness in terms of what we can understand as a sane, coherent reason, rather than just any bizarre reason that a psychotic person might respond to.

A complication: Unfortunately, there are problems with the present proposal for what counts as responsibility-generating control. The drunk driver example serves as an example of the trouble. As you may recall from a much earlier discussion in Chapter III, the trouble arises because the sufficiently drunk driver really does not have guidance control over the car or over their own behavior at the time of the accident. So, if control *at the time of the action* is required for responsibility, the very drunk driver is not responsible for the accident they cause. This seems wrong—we routinely blame drunk drivers. But, again as you may recall from Chapter III, there is a relatively obvious solution to this problem: the drunk driver lacks control at the time of the accident, but is nevertheless responsible because they had control over whether to get drunk in a situation in which they knew they would be likely to drive, and they knew the risks.

F & R propose to deal with the drunk driver problem in exactly this way.⁵¹ They agree that the very drunk driver does *not* have reasons-responsive processes guiding him or her

⁵¹Pages 49 -50.

at the time of the accident. So, to ensure that the drunk driver nevertheless counts as being responsible for the accident, they propose to alter their theory a bit to include events that result from choices that people have made earlier. Here is their final proposal:

When

- an action is guided by a person's own reasons-responsive choice-making processes, OR
- an event is a direct result of an earlier reasons-responsive choice-making process owned by that person and the person should have been able to foresee the possibility of that event's resulting from their choices at that time,

then that person has responsibility-generating control over that action or event.

The drunk driver counts as responsible, not because of the first item on the above list, but because of the second item. The first item on the above list is the proposal we had from F & R earlier on. The second item has been added to deal with cases like that of the drunk driver. Notice that according to the current proposal, the person has responsibility-generating control if *either one* of the two conditions on the list is satisfied.

Of course, this is not just about drunk driving. There are quite a number of situations in which people are reasonably thought to be responsible even though at the time they are not operating under the control of a reasons-responsive choice-making process.

It is important that the future results be foreseeable, and be foreseeable by the person who made the earlier choice. If I enter a booby-trapped building and thereby blow it up, killing dozens inside, but I had no way of foreseeing that it was booby-trapped, I am not responsible for the deaths, even though my choice to enter the building was made by a reasons-responsive choice-making process.

Animals and morally immune people. Look again at what it takes for a choice-making process to count as being reasons-responsive. Notice that some animals presumably can make choices that are reasons-responsive and that we certainly can understand as being coherent and sane. My dog, I am quite sure, qualifies. And yet, my dog is not capable of being morally responsible. He cannot understand moral reasons; he does not have a concept of moral duty or obligation; the notion of rights is foreign to him.

This really isn't a problem for the theory being presented as long as we see it only as a theory about responsibility-generating control and not as a complete theory about *moral* responsibility-generating control. I do not think there is a problem about saying my dog is responsible for certain choices that it makes, in the sense that these choices show what kind of personality or character it has. This talk of responsibility in animals would distinguish cases in which an animal's behavior was not under its own control from cases

in which the behavior is under the animal's control. However, what more needs to be added to the story about control in order to explain what it is to be *morally* responsible?

F & R have a proposal for us to consider to answer that question. The proposal is simple. *In order to be morally responsible, a person has to be able to recognize that sometimes "moral claims of others" constitute genuine reasons for him to act or to modify his actions.*⁵² This is why animals are presumably not ever morally responsible. They can't recognize the moral claims of others as reasons for themselves to act. Or at least that is what we commonly think about animals. If study of animal psychology revealed that some animals really do have a moral sense, an ability to recognize the moral claims of others as constituting reasons, then we should change our minds and see that those animals would be capable of being morally responsible agents.

What about people who are so immoral or anti-moral that no moral reason would ever actually move them? F & R would say that normally these people are still morally responsible for their immoral acts, because unless they are quite abnormal even these people *recognize* the moral claims of others—they *understand* other people's rights, for example—they just never allow those moral claims to have any real weight. They see that if they were to give weight to those moral claims, they would have to modify their actions. And, of course, they remain reasons-responsive—there is, we imagine, some potential that they could alter their actions if they were presented with sane and coherent reasons, even moral ones, since they recognize the force of moral claims —and so they have the right kind of control over their actions to be responsible for them. It is just that moral reasons are not among those reasons that they ever in fact allow to move them. So, moral responsibility, on F & R's account, does *not* require that people actually be moved by moral reasons—it only requires that a person be able to recognize the force of moral reasons.

This does mean that some mentally abnormal persons probably cannot be morally responsible for anything—at least this is the case if certain current theories about autism are correct. According to these theories, some autistic individuals cannot put themselves into the shoes of another person—they cannot understand another person's perspective at all. Such a person, we may imagine, cannot recognize the force of the moral claims of others, for they cannot understand what position others are in. So, if such an individual were to kick you hard in the leg, and then laugh at your pain-contorted face, they would not be morally responsible for their action, on the F & R account.

To me, this sounds reasonable. These individuals are reasons-responsive—that is they are controlling their own actions in a way that makes them responsible for those actions, but they are not *morally* responsible for them.

⁵²Page 79.

To sum up, we now have a two-part theory about the kind of control required for moral responsibility:

1. A person has responsibility-generating control over an action or event just in case—
 - the action is guided by a person’s own reasons-responsive choice-making processes, OR
 - the event is a direct result of an earlier reasons-responsive choice-making process owned by that person and the person should have been able to foresee the possibility of that event’s resulting from their choices at that time.
2. Additionally, in order to be *morally* responsible for the above event or action, a person has to be able to recognize how the moral claims of others can constitute genuine reasons for him to act or to modify his actions (but the person does not have to allow the moral claims of others to actually make a difference to the choice-making).

*The person’s own choice-making process.*⁵³ In the account given above, it is important that the choice-making process be the person’s “own.” There is more involved here than you might think. It is not enough that the process occur within the person rather than within some other person. For, a choice-making process might have been imposed on someone from the outside, in which case unless the person somehow makes the process their own, they cannot be morally responsible for the things that result from its application. Remember, F & R are trying to develop a theory of responsibility that works even though people are not ultimate moral authors—everything that goes on in a person may be caused by other things in the person or outside of them. So, it is important whether the choice-making for which a person is responsible “comes from” the person—that the choices are the person’s “own.” F & R need to have a careful test for when the process of choosing is one’s own.

According to F & R, as a person develops (presumably in childhood) they come to see themselves as someone whose choices make a difference to what happens around them, and that these choices bring about attitudes of moral approval or disapproval directed toward them. When, based on this sort of evidence, the developing child comes to accept the idea that it is fair for other people to react with approval and disapproval as a result of the way his choices are being made, the child has taken the crucial steps in “owning” his choice-making processes as morally relevant.

⁵³F & R discuss this issue in Chapter 8.

Whether or not the owning of a process occurs in childhood, or later, the elements involved are the same. People see themselves as choosers who affect the world around them by their choices, and they become familiar with the fact that they go through processes in making choices—sometimes thinking things through, sometimes going with the gut, sometimes trying to figure out what will feel the best, and so on. (These may well be unconscious in large measure—like programs that run in the brain.) But in order to “own” these processes, it is necessary for the person to accept the idea that it is all right for other people to evaluate them on the basis of the choices the processes produce. That is, when I employ these processes, I see them as “me” making the choices, not as some foreign force making the choices, and so I see it as fair that I be evaluated on the basis of the choices. If I see my choices as emanating from some force of power inside me that I do not identify with, that I do not see as “me,” then I will think it is unfair to evaluate me on the basis of those choices—since it was not “me” that made them.

This does not mean that a person has to consciously think this all through in order to “own” the choice-making processes, or that they have to think all the approvals and disapprovals they have received are fair. I certainly remember many times when I screamed at my parents about how unfair they were being. Rather, the idea is that I feel as though my choices are being *made by me* so that it is no wonder that people when they evaluate me use those choices as their basis.

So, this about an attitude I take toward the processes I use to make choices, even if I cannot say what those processes are. If I see my choices as emerging from somewhere other than “me,” I am not owning the processes used to make them. My moral identity will then be in doubt, and there will be trouble regarding just who is responsible for those choices. In the normal case, this problem will not arise, for usually people feel that almost all the time the choices they make are their own. There will be trouble about responsibility only in those odd cases in which the person who seems to have made the choice honestly does not feel as though the process was their own.

In mentioning those odd cases, I am not talking about situations in which the person feels threatened or pressured. In cases like that they may well not like the choice they feel forced to make. But they nevertheless recognize the choice as having been arrived at by use of their own processes. That sort of case is quite different from the case in which it truly feels to the person as though their choice-making process has been taken over by some alien force inside them, or by some external force that is acting on them, so that the choices that emerge are not made by a process that is their own.

Since F & R require that the choice-making process be one’s own before one can be morally responsible for its results, in those cases in which the person feels alienated from the choice-making process, F & R would say that the person is in fact not morally responsible. These are troubling cases, and I find it hard to see clearly what it might be reasonable to say about them. Cases of split personality come to mind as examples in which the position being taken by F & R seems reasonable. Such persons may not be morally responsible all the time. But what about cases like that of St. Paul in the quotation given earlier, wherein St. Paul laments the fact that he feels like some alien evil force inside him is making him choose the things he does not want to choose? I certainly don’t want

to say that St. Paul thereby loses his moral responsibility for those choices. So, I think there is trouble here for the F & R account. St. Paul feels alienated from his choice-making and yet do we really want to say that means he is not responsible for those choices? I don't think so.

The trouble arises when someone is of two minds about something important. That seems to me what happens with St. Paul. One part of him wants to be godly; another part of him wants the opposite. He may "identify" with the godly part, but the ungodly part still wins the contest over how he will live. If we go with the F & R proposal for what counts as a person's own choice-making process, I'm worried that St. Paul turns out not to be morally responsible for his evil choices, since he does not identify with the choice-making process that produces them. This seems to me to be an incorrect result that shows that F & R have not correctly identified what it means for a choice-making process to be one's own.

I think we should consider a different account of what it means for a choice-making process to be a person's own. Clearly, we do need to have a requirement that a person can be morally responsible only for choices that they make using a process that is in some sense their own. If they have truly been taken over by someone else, so that their choice-making really is not their own, then it would be unfair to hold them responsible for the results.

One easy "solution" would be to say that the choice-making process is one's own just in case one is completely self-determined. However, we already know people can't manage to pull that off. Realistically, people develop choice-making processes as a part of mentally maturing, but they do not generally deliberately develop them. Many of the details about choice-making are the result of experience, not created deliberately. Thus, the easy solution will not work, unless we want to destroy moral responsibility by requiring people to be ultimate self-determiners.

Here is a proposal that I think may be the proper solution to this open question: a choice-making process is your own when you came to use it through your own normal mental development, without morally excessive manipulation or interference, and without the presence of significant breakdown in sense of personal identity. It seems to me that this solves the St. Paul problem, since he does not lose his sense of personal identity—he always knows who he is, even though he doesn't always like what he sees. It also takes care of the need to exclude excessive manipulation and interference.⁵⁴ So, if this proposal works, whether or not the choice-making mechanisms you use count as your own will depend on the actual history of how you came to be using them. Whereas F & R put emphasis on the attitude you have toward the processes—whether you acknowledge them as your own, I would put emphasis on how you actually came to make choices in that way. I don't see that it matters whether you feel as though the processes are your own. After all, you might have been manipulated to feel that way, so your willingness to see your choices as coming from yourself may be the result of a clever plot. An unwitting pawn is still a

⁵⁴I think it is possible that my proposal is actually very close to what F & R meant to say. The St. Paul problem might be based on a misunderstanding of what F & R in the end want to say about a person's owning their choice-making mechanisms.

pawn. For this reason, I prefer my historical approach to the question of whether your choices are your own.

Summary. There is already a summary of the present theory found in the previous box, so I will not repeat that here. However, the discussion just completed defines what it means for a choice-making process to be one's "own." That definition needs to be applied to the material in the box. Take a look at it now. Really. Take a look before going on.

Again, this account of the sort of control required by moral responsibility does not see moral responsibility as needing to be based on contra-causal choice-making or ultimate moral authorship. It only requires using reasons-responsive processes of choice-making that are the result of normal human development, and a recognition on the part of the responsible person that moral reasons can be sufficient reasons for acting.

Critics will charge that this type of control is too weak to make the person genuinely responsible. They will say that since every step of the choice-making process might be caused by various things, the choices that result are not really those of the person involved.

Counterexamples. Unfortunately, even the modified theory presented above does not seem to be without flaw. Here is one troublesome example that the theory doesn't seem to be able to handle:⁵⁵ The driver whose mind wanders because of preoccupation and who thereby gets into an accident that injures someone is responsible for the accident even though the accident is not something they chose—they did not bring the accident about through the exercise of a reasons-responsive choice-making mechanism, and (more importantly) is not something foreseeable at the time of a prior reasons-responsive choice-making. True, the driver did choose to drive, but that choice is not the choice that brought about the accident. There does not seem to be any identifiable choice made by the driver that brought about the accident. The wandering mind of the driver is not something chosen by the driver, even though the driver could have foreseen the risk of bad consequences of a wandering mind. It is more like the driver fails to exercise the energy needed to keep her mind on driving. This is not like choosing to look away from the road to tune in the radio or make a cell phone call. It is more like mindless failure to exercise the required level of care. We would say the driver was negligent. This is not the same thing as choosing to do something other than paying attention.

Check back to the boxed summary of the theory. What in that box can make the driver morally responsible for the accident? I don't see that anything can do so. This is a problem. It is not just a problem about this one example. In general, there will be trouble about mindless negligence that immorally causes harm.

Here is another counterexample to the F & R theory: under many commonly occurring circumstances, the theory appears to blame the victim of a rape for being raped, even though she did nothing morally wrong.⁵⁶ If the victim was engaging in morally acceptable behavior, such as walking home through a dangerous neighborhood at night, when there was some foreseeable chance of being raped, then the victim made a choice that could be seen to carry significant risk of resulting in being raped. So, according to the F & R theory

⁵⁵I owe the example to Jodi Kane, one of my former students.

⁵⁶I owe this example to Jennifer Bode, one of my former students.

in the box above, she is morally responsible for being raped. Isn't this an unacceptable result; isn't this just another case of blaming the victim? And isn't the same unacceptable result going to be obtained for a whole range of cases in which crime victims could foresee some significant risk of becoming a crime victim?

It is not easy to see how to fix the theory to get around problems like these. If we simply remove the foreseeability clause from the theory, then the drunk driver gets off the hook. But if we leave that clause in the theory, we have trouble with blaming the crime victim. The clause apparently needs to be modified. But how?

I remain convinced this approach to moral responsibility is the most reasonable one I have talked about, but I am also convinced that it is in need of modification.

What do you think?