

EXCERPTS FROM JOHN MARTIN FISCHER'S DISCUSSION
WITH MEMBERS OF THE AUDIENCE

Scott MacDonald:

I was a bit surprised John about what you said about the case of Clara, who doesn't believe that she is morally responsible. You suggested that if one lacks the relevant belief, or has the relevant negative belief, one will be like the pilot of the ship who lets the boat be buffeted by the winds. Why should you be inclined in that direction given what you said in reply to Carl [Ginet] earlier? You claimed that knowing that there is bound to be a single outcome does not mean that we would stop deliberating because we don't know what it is. Of course, that is a standard compatibilist line about that sort of thing. I would have thought that exactly the same sorts of considerations would apply here. Just because you don't think that you are responsible doesn't mean that you would stop engaging in those things. You would just cease having certain beliefs about how people should react to you once you have done all of that.

John Martin Fischer:

Yes. I think that I agree. In other words, I think that you would still go on deliberating. I think there was something perhaps misleading about saying that you are just like the sailor who allows the boat to be buffeted. But remember the addict (keep in mind that Frankfurt uses the term "addict" for a person who knows that he has an irresistible urge to take a drug). Such an agent may go ahead and deliberate about how exactly to get the drug and then proceed to do so. But all along the addict thinks that he is an addict and therefore it would be futile even to try to avoid taking it. Frankfurt's intuition is that the agent is not responsible (as long as he is not responsible for becoming an addict in the first place). So I think you could still deliberate and so forth. But if you don't see yourself as responsible, I don't believe you genuinely can be responsible. Now, let me admit, this is definitely one of the intuitively jarring parts of the theory. That is why I said that I would hope that people would judge it by its consequences as a whole: whether it is illuminating, it solves certain puzzles about actions and omissions and under what conditions we hold people responsible for



the consequences of what we do. Also does the theory deal with various arguments for incompatibilism?

Let me say one more thing. It is not as though "anything goes," if one is not morally responsible. Someone who is not responsible can still be sequestered, can still be treated as if he were an animal, or insane, or simply treated harshly. It is just that it would not be *punishment*, or it would not be *full moral responsibility*.

Carl Ginet:

My initial resistance to your subjective view is that someone could simply avoid blame and all things that go with it by endorsing certain beliefs (of a hard determinist sort). But in your defense, they cannot just make themselves believe that they are not responsible.

John Martin Fischer:

Right! That's the problem with that sort of objection to my view. You cannot just say, "By the way, I don't believe that I am responsible." When you see how minimal the conditions are, it is clear that it is only in the rare case that they will go unsatisfied. Remember, on my view, the conditions are just a matter of seeing that one is causally efficacious in the world, that one is an apt target for the reactive attitudes given the social practices of the community, and that one bases one's evidence for these beliefs in the appropriate way.

Carl Ginet:

It would seem that, if believing that you are not responsible (say because you are an incompatibilist and you believe that causal determinism is true) does not lead to drastic differences in your behavior, your psychē, your whole thought processes about your actions, then you don't really believe it. You are just saying it. You don't have the dispositions that would be expected with such a belief. Now a person who *really* believed that he is not responsible would be a strange, strange creature. We might want to say that he was crazy. And we then might not hold him responsible.

John Martin Fischer:

I appreciate the help. But I may be in a little more trouble. In other words, it may be that someone could at least sincerely doubt whether he is an apt target of the reactive attitudes in a deterministic world. That is obviously not so crazy. But what I would want to do is distinguish a reflective versus a non-reflective version of the condition. But, you are right. And I would definitely emphasize that it is not just a matter of what you say. It is a matter of what you genuinely believe. Even if you genuinely believe that

you are not responsible you still could be sequestered from society. You still could be treated harshly. It is just that you would not be participating in the conversation that is responsibility.

Joseph Margolis:

I just need a little clarification I think. As you were speaking I found myself thinking that some of these questions come out a little more clearly when one reflects on the context of legal cases. I was trying to think of how restricting discussion to moral cases, one could make out what the proper sense is in which the relevant terms are used. My intuition is that, in the moral cases, we usually use the term "responsible" as an indication of blame. If someone is blameable for an action then the usual sense is that he is responsible for it. There may be more to it. There may be other sorts of cases. For example, in the legal case, parents are often held responsible for what their children may do. But they are not necessarily blameable for it, and in that sense they may not be responsible for what their children have done. The usual argument is, "I don't think that I should be held responsible because I am *not* responsible." We collapse the two senses in a certain way. That suggests to me the following: it just will not do either to reduce concrete cases to abstract principles or to appeal to wild scenarios. That is, we begin with are the garden variety cases of persons being held responsible. And then we expect that if someone wants to get out of the charge, then he or she will mention extenuating considerations, also garden variety, in virtue of which one will say, "Oh I didn't realize that." But there is no rule for the kind of extenuating circumstances you want us to consider. All the problems that seem to generate difficulty have to do with speculation about a family of extraordinary, not garden variety sorts of cases.

John Martin Fischer:

Good. In our session this morning I suggested that there are different routes to the conclusion that we do not need alternative possibilities for responsibility. One route is taken by Jay Wallace and I think that it is very similar to what you are saying. It is looking very carefully at our practices of praising and blaming and exculpation and trying to figure out what is encoded in those practices. His argument is that if you look carefully there is no underlying principle of alternative possibilities there. That is, I think, extremely useful and I absolutely welcome that sort of analysis; in Wallace's case, it issues in the same conclusion that I draw. But in philosophy often it can be helpful to have different routes to the same conclusion. I agree that sometimes this literature on the Frankfurt cases focuses on somewhat wild

examples, and one should never rest one's case *solely* on such examples. One thing I would say is that there is a great tradition in philosophy of thought experiments, perhaps going back to Gyge's ring and Descartes' evil genius. My approach is in that tradition.

One last thing: I often think that these scenarios aren't so wild. For instance, this really did happen. I went to the grocery store quite some time ago and the person behind the cash register said "Paper or plastic?" I said, "Well, paper is fine." She looked down and she said, "Oh, good thing. We don't have plastic anyway." I thought, now wait a second, I just lived through a Frankfurt example! So I don't know that it is so wild.

Joseph Margolis:

Just to press one step further: Your view about believing that you are responsible seems to be derived from a very high order theory rather than from the garden variety cases. So there is some question as to why we should bring that in. It wouldn't be obvious in that way.

John Martin Fischer:

Well, OK. Let me say a couple of things. It turns out that Bob Kane has a section of his book in which he talks about moral education and coming to be a moral agent. This is very similar to the kind of approach that Mark Ravizza and I develop. The conditions we have on taking responsibility actually do *not* get motivated on our view by high-minded theory or complex and arcane examples, but rather by reflecting on how kids become moral agents. Now, admittedly, it is a bit of arm-chair moral developmental theory. But we basically want to say that as children come to see themselves as agents, and come to be morally educated, there are typically certain stages. Our account is supposed to mirror that. By the way, our account of taking responsibility is there to capture what goes on in moral education. It is then a serendipitous implication that it also helps with the manipulation cases.

Mark Case:

Your phrase "taking responsibility" reminds me of another phrase that sometimes gets used in these kinds of contexts, which is "owning up to your action." It seems like we have the event which is the action, and then there is the coming to believe that one is responsible. There are a couple of cases that come to mind. One is that you do the action and at the time that you do the action you believe that you are not responsible for the action. You believe that the addiction is fully efficacious. Then later on you come to think on the basis of some compelling evidence that you were

responsible. Then there is the reverse kind of case where, at the time which you are doing the action, you believe that you are responsible but later you come to believe otherwise. I was just wandering what you would say about these kinds of cases.

John Martin Fischer:

First, I liked your point that there is another phrase we use and that is "owning up to your action." Now maybe it is just a linguistic accident, but I think of taking responsibility as acquiring a certain kind of ownership which is the same word that we use when we say "own up" (of course this could just be a linguistic accident, and I do not want to adopt a Derridian methodology!). But I would say that if you don't look at yourself as meeting certain minimal conditions, then you don't see yourself as responsible when you are acting. If so, then you are not responsible for your action even if later you come to change your mind. Secondly, consider the case where you do see yourself as responsible and later you change your mind. You might well be responsible.

Let me back up for a second. Mark Ravizza and I try to give an account of the different items for which we assign responsibility: actions, omissions, consequences, and emotions and character traits. We try to show that they fit together in a certain systematic way. The idea is that there are two main components to the analysis. There is the ownership or "taking responsibility" component. But also there is the reasons-responsiveness component. Let us say that someone acts on a mechanism that is not reasons-responsive and later comes to believe that he was in fact morally responsible. He might just be wrong. But suppose he does believe that he is an agent and an apt candidate for the reactive attitudes and so forth, and his mechanism of action is reasons-responsive, and he later says "Oh, gosh, in reflecting on this I don't really think that I was responsible." Well, tough luck, as far as I am concerned. He was responsible and later he comes to a different (incorrect) judgment.

Elie Noujain:

I think that Professor Kane will agree with you about the historicity of responsibility. But where I think that he will not agree is with your further point that free agent's history could be fully determined. He seems to think that moral deliberation is necessary but not sufficient. He has the added condition of *causal indeterminism*. You think that moral deliberation is sufficient. My question is, do you think that moral deliberation is sufficient *regardless* of whether determinism is true? Or do you think that causal determinism *must* be true?

John Martin Fischer:

My position is that it does not matter whether determinism is true *or not*. One of my motivating engines in trying to argue that determinism is compatible with responsibility is that I think that our view of ourselves as persons and as morally responsible agents should not hang by a thread. It should not be held hostage to the possibility that some consortium of scientists will discover and announce tomorrow that causal determinism is true. If they announced that, I would not stop holding people responsible. Of course, initially there would be doubt. But if after years we actually agree, if all of the physicists said yes, Bohm's interpretation of quantum mechanics is the right one and determinism is true, I would not stop looking at my daughter as deeply different from the kitten. I just would not give up my notion of moral responsibility. I want to capture that. It would be odd, wouldn't it, if I said, "Yes, but our attitudes towards other persons and our belief that we are morally responsible does hang on a thread in the sense that, if the scientist discovered that *indeterminism* is true, then we would have to give up our attitudes"? That is why I call myself a super-compatibilist. I don't wear a cape or anything. But I am a super-compatibilist in that I think that responsibility is compatible with determinism and with certain kinds of indeterminism, as well.

Let me say one more thing. You might think that the doctrine super-compatibilism (i.e., the doctrine that determinism is compatible with moral responsibility and indeterminism is compatible with moral responsibility) must be very far apart from the doctrine that Derk Pereboom holds, which he calls "hard incompatibilism." Hard incompatibilism for him is the view that determinism rules out moral responsibility *and* indeterminism rules out moral responsibility as well. What two views could be more different? Well, what Derk and I agreed on last night, maybe it was because we had a couple glasses of wine, is that our views are actually very close. My view is that determinism could be true and certain kinds of indeterminism could be true and we could still have what we really care about, which is, in my view, responsibility (now on my view, if determinism is really true, we would have to give up alternative possibilities, but the point is we would keep moral responsibility – the thing we really care about). Derk's view is that if determinism is true we have to give up something – i.e., the robust reactive attitudes that we thought we had – but we can still have what we really care about, something like the reactive attitudes, but not quite the robust version. Similarly for indeterminism. Thus, semi-compatibilism and hard-determinism are startlingly similar. One might have thought that the doctrines of semi-compatibilism and hard determinism were at opposite ends of the spectrum; but perhaps it turns out that the doctrines lie on a

Moebius strip, or one of those M.C. Escher staircases that turns back on itself!

Robert Kane:

By the way, I would chime in that I agree with both of you. We would have to give up the robust versions of freedom and moral responsibility (i.e., ultimate responsibility) if determinism were true. But surely I would go on living like everybody else with the non-robust versions. I would just feel like I had lost something.

Joseph Margolis:

Why wouldn't that be equivalent to saying that it is irrelevant?

Robert Kane:

Well, I would figure, Joe, that I had lost something. I would be living in a world that is less rich and valuable.

John Martin Fischer:

What I would say is that we would be living in a world that is different from what we thought, perhaps less rich, but not less valuable.

Carl Ginet:

I would feel, if I went on living like I did before I learned that determinism is true, that I was living an illusion, though I could not help doing so but that I was. It would be deeply disturbing the whole time. I am disturbed enough as it is.

I would like to lay on John a wild example, maybe the ultimate manipulation scenario. In your treatment of the Walden Two case you say, "Well I have hope that we can find some relevant difference that will strike us all, some account of the difference." I was thinking about your response to that kind of example and I thought, well, here is the example we really need: Suppose there are some super-human intelligences somewhere in the galaxy and they come to understand the laws of nature thoroughly. Suppose that determinism is true. At least they understand enough about the laws of nature that when they look at what happens on our planet they can understand exactly why everything happens. They can explain it causally. So one of them decides to run a sort of super *Truman Show*. What they do is they take an actual human life and observe it from the beginning. Then they say, we are going to reproduce that life in all its details. And they do. They know in advance how to do it. They are super human intelligences and can hold in their minds all of the details. They create the same embryo, etc., the same environment of this actual normal person. It could have been

you, or me. Now here, I take it that there are no relevant differences to the history of the person or their surroundings. I myself have a clear intuition that in this Super Truman Show the person is not morally responsible for any of their actions.

John Martin Fischer:

Yes, that is where we disagree. I think that it shows why you are naturally inclined toward incompatibilism and I toward compatibilism. My intuition is that world you described might be the actual world, if there is a God and God has certain kinds of knowledge and set it up to work in a certain way. My own view is that it might have been the case that God set up the world and knew in advance its entire history, maybe through middle knowledge or some other mechanism. But as long as what we do issues from reasons-responsive mechanisms, it does not eliminate our responsibility.

Maybe this is one way of understanding the issue. I think that there is an interesting difference between God on the occasionalist picture and on a different picture. On the occasionalist picture, God is continually intervening in the human mind. Imagine, let us say, a scientist, or super-human intelligences that come to us from another planet. In one case – as in occasionalism – they get us to do everything they want us to do by directly stimulating our brain at every moment such that that type of stimulation renders the agent not responsive to reason. On another type of scenario they just start the world in such a way that they know it will play itself out as they wish it to. I think that there is a crucial difference between the occasionalist picture and this latter picture.

Derk Pereboom:

Earlier you said that in manipulation cases, the agent does not take responsibility for the manipulated mechanisms. But are you willing to say now that if the manipulation is global enough or remote enough, then the agent does take responsibility for the manipulated mechanisms?

John Martin Fischer:

I distinguish two kinds of cases. One is where you take a baby before the baby becomes a moral agent. Scientists come and start manipulating the baby's brain. I would say that that baby never becomes a person, because the baby never develops into a moral agent. However, I would distinguish that kind of case from the case in which an individual grows up in the normal way (whatever that is) and takes responsibility for the mechanism of practical reason. Then, at some point, the scientists come in and start

manipulating the brain. In that case you have a moral agent who simply is not responsible for the manipulated behavior.

So I would distinguish those two kinds of cases. If you ask, if the manipulation is global enough, do you take responsibility? Well, if it is global in the sense that the baby never becomes a moral agent, then I would say that you don't even have a person there at all.

Derk Pereboom:

I see. So, what about this kind of case. You take the *Truman Show* thing, and suppose that in this case the manipulator manipulates constantly, but only at the micro-physical level. Still, the neural states and psychological states are indistinguishable from those of those of the normal person. What would you say about taking responsibility for mechanisms in that situation?

John Martin Fischer:

So in this case, someone is being manipulated, but the manipulation is indistinguishable at the neural level and at the psychological level from "ordinary" human functioning. This is an interesting question. In this case there is no reason to say that the mechanism is not reasons-responsive (insofar as the ordinary brain is). But I would still say that it is not the agent's own, unless he knows about the manipulation and consents to it.

This seems to be the hard case: from the very beginning someone is being manipulated, but the manipulation creates processes that are just like ordinary brain processes. I'd say that the individual never becomes a person.

Derk Pereboom:

Why would these manipulated mechanisms not be the agent's own? You want to allow that normal, non-manipulated agents might take responsibility for mechanisms they know little about, or about which they have false beliefs. If, for example, an agent can take responsibility for his psychological mechanisms not knowing that they have neural constitutions, or believing that they do not have neural constitutions, why can't my manipulated agents take responsibility for their psychological mechanisms not knowing that these mechanisms have constitutions that are manipulated at the microphysical level, or believing that these constitutions are not manipulated in this way?

John Martin Fischer:

I think that this is a good challenge. I think that I am committed to the view that manipulation that is significant from the very beginning – that does not allow the individual ever to morally develop into his or her own person – *does* rule out responsibility. This is so even though the brain processes might be type-identical to those that go on in someone who is not being manipulated.

Derk Pereboom:

But what makes it the case that the manipulated person does not develop into a person?

John Martin Fischer:

It is because the mechanisms on which he acts are not his own.

Robert Kane:

Let me give you another variation on this (I know that we are getting more and more bizarre). Suppose that some future scientists discover a way to do a kind of incubator version of raising a child to be ten years old. Instead of having to go through all of that difficult stuff that most of us here went through of raising a child to be ten years old and to be a certain kind of moral agent, the scientists say, "Hand us the child. We will put it in the incubator and six months later we will duplicate the process and give you a ten year old who doesn't steal cookies and is otherwise fairly well behaved."

John Martin Fischer:

I think that one can have, let us say, just artificially produced individuals that are molecule for molecule isomorphic to us and yet they are not responsible because of their history – until they have taken responsibility.

Robert Kane:

But it seems that both of these ten year olds can take responsibility in your sense.

John Martin Fischer:

Well, they might say that. But I would want to make sure that the relevant conditions are met. For instance, I think that there can be instant agency but not instant responsibility.