


This article was downloaded by: [University of Notre Dame]

On: 24 February 2011

Access details: Access Details: [subscription number 731689697]

Publisher Routledge

Informa Ltd Registered in England and Wales Registered Number: 1072954 Registered office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH, UK



Philosophical Explorations

Publication details, including instructions for authors and subscription information:

<http://www.informaworld.com/smpp/title~content=t713706422>

Reply: The free will revolution

John Martin Fischer^a

^a Department of Philosophy, University of California, Riverside, California, USA

To cite this Article Fischer, John Martin(2005) 'Reply: The free will revolution', Philosophical Explorations, 8: 2, 145 – 156

To link to this Article: DOI: 10.1080/13869790500091516

URL: <http://dx.doi.org/10.1080/13869790500091516>

PLEASE SCROLL DOWN FOR ARTICLE

Full terms and conditions of use: <http://www.informaworld.com/terms-and-conditions-of-access.pdf>

This article may be used for research, teaching and private study purposes. Any substantial or systematic reproduction, re-distribution, re-selling, loan or sub-licensing, systematic supply or distribution in any form to anyone is expressly forbidden.

The publisher does not give any warranty express or implied or make any representation that the contents will be complete or accurate or up to date. The accuracy of any instructions, formulae and drug doses should be independently verified with primary sources. The publisher shall not be liable for any loss, actions, claims, proceedings, demand or costs or damages whatsoever or howsoever caused arising directly or indirectly in connection with or arising out of the use of this material.

REPLY: THE FREE WILL REVOLUTION

John Martin Fischer

Introduction

I should like to thank the editors of *Philosophical Explorations* for kindly offering an opportunity for discussion of some of my (and my co-author's) work on free will and moral responsibility. Additionally, I am very grateful to Ishtiyaque Haji for his excellent work in facilitating this symposium and editing the papers.

My deepest debt of gratitude goes to Michael McKenna, Daniel Speak, Neal Judisch, and Seth Shabo for their probing, gracious, and illuminating discussions. I am very thankful for their thoughtful and sympathetic papers, from which I have learned a great deal. I shall limit myself to some sketchy and tentative replies, and I hope to continue to think about and work on these issues in the future. I cannot possibly do justice to all the points raised in these papers.

The Structure of Semicompatibilism

I have argued that moral responsibility does not require the sort of freedom or control that involves genuine metaphysical access to alternative possibilities ('regulative control'). Rather, 'guidance control' is the freedom-relevant condition necessary and sufficient for moral responsibility. When an individual controls his behavior in the pertinent way, he acts from a certain sort of general capacity for reasons-responsiveness. A bit more specifically, an individual exhibits guidance control of his behavior insofar as it issues from his own, moderately reasons-responsive mechanism.¹ On this sort of view, an agent need not have access to alternative paths into the future, in order to be morally responsible; the relevance of what happens along non-actual pathways is to specify the 'modal' or 'dispositional' features of the general capacity in question. Semicompatibilism holds that causal determinism is compatible with moral responsibility, quite apart from the issue of whether causal determinism is compatible with genuine metaphysical access to alternative possibilities. The future may—or may not—be a garden of forking paths (in Borges' notable phrase); nevertheless, we can legitimately be held morally responsible for how we take our path into the future.²

I have contended that it is a significant virtue of semicompatibilism that it allows us to sidestep seemingly irresolvable disputes about the relationship between causal determinism and regulative control. On my approach, possible scenarios with different laws from the actual world and/or different pasts from the actual world are indeed relevant to the question of moral responsibility; as I said above, they are relevant in virtue of helping to fix the modal characteristics of the actually operative mechanism (the actual-sequence

mechanism that issues in the behavior under scrutiny). But it can be puzzling how exactly I think progress can be made by turning away from the issue of whether we have regulative control but at the same time focusing on non-actual possible worlds (or scenarios) in which the past and/or laws of nature are allowed to differ from actual world—precisely the sorts of scenarios typically shunned by incompatibilists.

Both McKenna and Speak press this important worry about the structure of semicompatibilism. McKenna says:

Given the nature of determinism, there is no possible world which shares the same facts with the actual world at a time and the same laws but in which an agent acts differently than she does in the actual world.

It appears that Fischer and Ravizza's reply to the above incompatibilist's objection *must be* that what matters for assessments of moral responsibility is the *general capacity* to react to reasons, and not the particular ability of an agent to act differently to a certain sort of reason. And this is, indeed, what Fischer and Ravizza say . . . Unfortunately, they offer little argument for this crucial contention. They maintain that it is plausible to reply to the incompatibilist that an agent's general capacity to react to different reasons to do otherwise in other possible worlds shows that an agent 'could in fact have reacted' (i.e., possesses the general capacity to react differently) to her actual reasons for doing otherwise . . . But the plausibility of Fischer and Ravizza's reply will be lost on the incompatibilist. No doubt, the incompatibilist will acknowledge that the pertinent worlds do demonstrate that such an agent does possess the *general capacity* at issue. But Fischer and Ravizza will need to say more about why the possession of such a general capacity should satisfy the incompatibilist at this point. Why exactly does shifting to properties of the mechanism (or general abilities of the agent) help here? In the absence of an answer to this question, the apparent, long-standing philosophical stalemate remains. (McKenna 2005)³

Similarly, Dan Speak says:

. . . on Fischer and Ravizza's account, a person can properly be held morally responsible only if there is another possible world with a different past at which the agent does otherwise.

This begins to bring out my concern. After all, one way to conceive of the incompatibilist's central intuition is that moral responsibility can only be assessed with regard to the *actual* past. Who cares, the incompatibilist will ask, what the agent would have done under different conditions? Even the most basic machines will produce different outputs when the inputs are varied. To carry moral weight, reactivity needs to have more than this merely counterfactual basis.

This might, in fact, be the rationale behind the standard objection to the conditional analysis of freedom employed by traditional compatibilists . . . the conditional analysis fails in large part because it appeals to possible worlds that are inaccessible to the agent in question. There are cases, as the brainwashed person shows, where we think of freedom in terms of what the agent could do under the conditions in which he found himself, regardless of what might have happened had the past been different. Now the incompatibilist will, plausibly, want to say the same thing about the sort of reactivity to reasons that is relevant to moral assessment.

Why is this important? Because Fischer and Ravizza's semi-compatibilism was supposed to help break a deadlock—offer another move in the face of the dialectical stalemate that is

the free will debate. Now that we understand the constitution of reactivity, we can see that there is reason to withhold confidence in Fischer and Ravizza's ability to advance the debate in the way they have hoped. One central attraction to adopting semi-compatibilism was that in doing so we might be able to by-pass the older debate about which counterfactuals were relevant to the determination of powers and abilities. But if Fischer and Ravizza are going to respond to the incompatibilist objection to weak reactivity that I have outlined, then they will have to enter this debate after all . . . [In doing so] they will be re-inscribing the traditional debate, and, therefore, will not be delivering on their promise of being able to execute and end-around to avoid the historic modal arguments. (Speak 2005)

To reply: begin by distinguishing three different approaches to moral responsibility. The first requires alternative possibilities, or regulative control, for moral responsibility. On this approach, the individual whose moral responsibility is at issue must have been free to do otherwise—must have had genuine access to another possible scenario—at some (relevant) point along the path to the behavior in question. This approach to moral responsibility could be adopted by either a libertarian about causal determinism and regulative control or a compatibilist about causal determinism and such control. So, for example, this approach to moral responsibility is embraced by both Peter Van Inwagen (a libertarian) and also Keith Lehrer (a compatibilist).⁴

On another approach, moral responsibility requires the mere presence of a general capacity for reasons-sensitivity, *not* regulative control (at *any*) point along the pathway to the behavior under consideration. Note that, on this approach, an agent may be morally responsible simply in virtue of possessing the relevant general capacity; he need not have exercised that capacity in the circumstances. A proponent of this sort of approach to moral responsibility is R. Jay Wallace (1994).

My view is, as it were, 'in between.' That is, on my approach, regulative control is not required at any point along the pathway to the behavior. But the mere presence of the relevant general capacity is *not* sufficient. The freedom-relevant condition necessary and sufficient for moral responsibility, on my approach, is that the behavior *actually issue from* the relevant general capacity: the agent's own, appropriately reasons-responsive mechanism. So I would take issue with McKenna's formulation, according to which 'what matters for assessments of moral responsibility is the general capacity to react to reasons.' Better: I would urge that what matters is not the *mere presence* of such a capacity (as in Wallace), but the actual *exercise* of the capacity.

Now distinguish two versions of the actual-sequence approach to moral responsibility I favor (according to which moral responsibility does not require the sort of control that involves genuine metaphysical access to alternative possibilities, but does require that the actual-sequence mechanism issuing in the behavior have certain characteristics). One version has it that those characteristics are entirely non-modal (or non-dispositional), whereas the second version allows the relevant features to be modal or dispositional. I prefer the second version. Given that one can make sense of modal or dispositional properties (at least in part) in terms of non-actual possible worlds, on my favored approach to moral responsibility non-actual possible worlds (or hypothetical scenarios of certain sorts) are pertinent to moral responsibility, but *not* in virtue of indicating genuine metaphysical access; rather, they help to specify the sorts of properties that must be exemplified by the actual-sequence mechanism that issues in the behavior in question.

But why exactly does this 'shift' in the point or purpose of the assessment of non-actual possible worlds help to make dialectical progress? As Speak and McKenna point out, the incompatibilist typically insists that the only non-actual possible worlds relevant to assessments of moral responsibility are the ones that share the past and natural laws with the actual world. Why care at all about worlds in which the past and/or natural laws are different?

The answer is that the *purpose* of the assessment of alternative scenarios makes the crucial difference. If the point or purpose of the invocation of alternative scenarios is to establish that one has genuine metaphysical access to alternative possibilities, then I grant that only such scenarios with the same past and laws as the actual world are relevant; this is because our freedom is the freedom to extend the given past, holding fixed the natural laws (Fischer 1994; Ginet 1990). But this is not the purpose or 'use' to which I put alternative scenarios in my approach to moral responsibility; rather, the purpose of the assessment of ranges of non-actual possibilities is to establish the nature of the properties that are in fact (actually) exemplified in the sequence of events leading to the behavior under consideration.

I contend that, given this difference in purpose or employment of the apparatus of possible worlds in my theory (as contrasted with theories that posit the requirement of regulative control for moral responsibility), it can be seen that it is *not* problematic to invoke non-actual possible worlds with different pasts and/or natural laws than those of the actual world. One way to see this point is to recognize that the problems that are raised by the invocation of such scenarios in the context of regulative control *do no arise* in the context of ascertaining the modal properties of the actual-sequence mechanism.

To explain: the general problem with the 'conditional analysis' of freedom is that some 'obstacle' to one's doing *X* may actually exist and render it true (intuitively) that one cannot (in the relevant sense) do *X*. Now it may be true that if one were to (say) choose to do *X*, then the obstacle would (for whatever reason) be gone, and thus the following counterfactual may be true: if one were to choose to do *X*, one would do *X*. So the relevant counterfactual conditional is evidently true, and yet intuitively one cannot do *X*. The range of possible worlds in which the actually-existing obstacle is 'subtracted' is not relevant to the assessment of the issue of whether one possesses the ability in question in the actual world. Similarly, since an individual's freedom is the power to extend the actual past, holding fixed the natural laws, the range of possible worlds in which the past and/or the natural laws are different from those in the actual world are (arguably) irrelevant to the assessment of one's ability (in the relevant sense). Shifting to such worlds is analogous to subtracting some actually existing feature that intuitively made it the case that the individual lacked the ability in question (in the actual world). The non-actual possible worlds are crucially different from the actual world in some respect that renders them irrelevant.⁵

But nothing like this is going on, when one invokes non-actual possible worlds (hypothetical scenarios) in seeking to give an account of the modal or dispositional characteristics that are actually displayed in the sequence leading to the behavior. We are not here entertaining the issue of access to non-actual possible worlds, which may in fact be blocked by some feature of the actual world (or the relationship between the actual world and accessible non-actual possible worlds). We are simply giving an account of the modal characteristics of the actual sequence of events. Thus, the specific problems that arise when considering the conditional analysis of ability to do *X* (or, for that matter, other compatibilist accounts of the relevant sort of ability) *do not* come up in this context.

To help to see my point, consider the following example. Incompatibilists will point out that, no matter how well one can play the piano, one cannot (in the sense that is connected with moral responsibility) play piano in a particular set of circumstances in which one is chained to one's desk. One here may have the general capacity to play the piano, but one cannot exercise this capacity, and thus one cannot, in Austin's 'all-in' sense of 'can', play the piano (in the given, particular circumstances). Given the context, playing the piano would require the 'subtraction' of an actually-existing obstacle to playing the piano. But of course it does not follow from this fact that the individual does not actually have the general capacity to play the piano.⁶

Similarly, it may well be true that the only way a given agent could have reacted differently to an actually existent reason to do *X* would have been for the past or the laws to have been different. It follows, in my view, that the agent does not have access to a non-actual possible world in which he reacts differently to the actually-existing reason to do *X*. But, given that there are certain worlds in which the past and/or laws of nature are different in certain respects and the agent reacts differently to the actually-existing reason, he may well be acting from a general capacity that has the property of reasons-responsiveness. The pianist cannot play in the circumstances, but he has the general capacity to play the piano; the agent cannot (in the particular circumstances) react differently to the actually-existing reason, but he may well act from a reasons-responsive mechanism.

Note that it is clearly inappropriate to continue to use all of the conceptual and theoretical framework of a 'paradigm' or scientific worldview that has been superseded, or perhaps to apply this apparatus to a domain to which it has been shown to be inappropriate. So it may well be inappropriate to apply (parts of) the conceptual framework of Newtonian mechanics to certain quantum phenomena, or the axioms of Euclidean geometry to non-Euclidean or Riemannian spaces, or the postulates of String Theory (with its ten dimensions) to Membrane Theory (with its eleven dimensions), and so forth. Similarly, it is inappropriate to continue to insist that only those worlds which share the past and natural laws with the actual world are relevant to an evaluation of moral responsibility, once one has shifted from the paradigm of regulative control to the paradigm of 'guidance control.' This would be like keeping the idea of an absolute spatio-temporal framework while adopting the other postulates of relativistic physics.

The shift from a requirement of regulative control to a requirement of guidance control does not in itself provide a decisive defense of compatibilism about causal determinism and moral responsibility. But it does shift the debate into a dialectical terrain considerably more hospitable to compatibilism.⁷ If we do not need genuine metaphysical access to non-actual possibilities, a compatibilist need not deny the fixity of the past or the fixity of the natural laws. The semicompatibilist conceives of the role of hypothetical scenarios in a distinctive way; thus, the theoretical role of these scenarios is fundamentally different in traditional compatibilism and semicompatibilism. To mix elements of the older ways of thinking with the new only breeds confusion.

I proclaim, then, with all due modesty: Semicompatibilism is the Free Will Revolution. Agency theorists of the world, unite!

Mechanism Individuation

Semicompatibilism has it that the freedom-relevant condition necessary and sufficient for moral responsibility is guidance control. On my view, guidance control has two

distinguishable elements: mechanism ownership and reasons-responsiveness. One exhibits guidance control insofar as one's behavior results from one's own, moderately reasons-responsive mechanism. I (and my co-author) have sought to give accounts of each of these components.⁸

My approach to mechanism ownership is 'subjective'; one makes a kind of mechanism one's own by *taking responsibility* for acting from that kind of mechanism. On this sort of view, one acquires control at least in part by *taking control*. I (and my co-author) offered the following three conditions for taking responsibility:

- (a) The individual must see himself as an agent; he must see that his choices and actions are efficacious in the world.
- (b) The individual must accept that he is a fair target of the reactive attitudes as a result of how he exercises this agency in certain contexts, and
- (c) The individual's view of himself as specified in the first two conditions must be based, in an appropriate way, on the evidence. (Fischer and Ravizza 1994, 210–13)

Further, on our view, one takes responsibility for acting from certain general kinds of mechanisms. (An individual is morally responsible for acting from a particular mechanism only if he has taken responsibility for acting from mechanisms of that kind.) Employing this sort of account, I and my co-author suggested that we could say that someone who is being manipulated in a covert way (without his consent) has not taken responsibility for the kind of mechanism that issues in his behavior, and thus does not own the mechanism. Such an individual does not exhibit guidance control. Even if he has taken responsibility for ordinary practical reasoning, he has not thereby taken responsibility for the quite different kind of mechanism—manipulation by scientists.

Seth Shabo raises an interesting worry about the three conditions for taking responsibility:

... (c) *seems* to allow that the requisite self-conception could arise in a proscribed way even if nothing in the agent's history renders it non-veridical. For this reason, talk of evidence appears misleading. Evidence of *what*? The obvious answer—evidence *that* the agent's self-conception is veridical—seems to be precluded. For since the agent's *being* a fair target of the reactive attitudes is partly constituted by his so viewing himself, there is no independent reality against which to gauge the accuracy of this view. (Shabo 2005)

I am worried about the problem expressed by the sentence, 'For since the agent's *being* a fair target of the reactive attitudes is partly constituted by his so viewing himself, there is no independent reality against which to gauge the accuracy of this view.' But I am not convinced that the account is circular or problematic in the suggested way. I have offered a number of conditions for taking responsibility, including (c), where (c) is in part subjective. It posits that in order to be morally responsible, I must believe I am, and that that belief must be based on my evidence for the proposition (in the right way). Whatever it is to be an agent and a fair target for the reactive attitudes, I surely can have evidence for this. Further, I can have evidence that I believe I am an agent and am a fair target of the reactive attitudes—I can presumably have introspective evidence that I so believe.

Analogously, someone might say (à la Galen Strawson) that an agent *S* is free if and only if *S* meets certain (non-subjective) conditions plus *S* believes he is free (Strawson 1986).

There does not appear to be a circularity here, as the freedom claim is embedded in an epistemic (belief) context. Further, if one added some sort of evidence-basing requirement, this would not change matters.

Both Shabo and Neal Judisch press important and worrisome questions about mechanism-individuation. Mechanism-individuation is required for both components of the account of guidance control—mechanism-ownership and reasons-responsiveness. I have discussed these matters in some detail (although, no doubt, not sufficient detail) elsewhere (Fischer 2004a), and I shall only touch on them briefly here. I fully admit that I do not have a reductive account of mechanism-individuation, and that, lacking such an account, the theory of guidance control is, at best, incomplete. But an incomplete theory is not thereby an inadequate or fundamentally misguided theory. And my approach to guidance control and moral responsibility is no different from most promising philosophical theories in being incomplete. Further, some important philosophical theories are incomplete in a similar respect; theories such as reliabilism in epistemology and generalization in ethics (of both the consequentialist and non-consequentialist sorts) rely implicitly on views about the individuation of the relevant kinds of belief-producing mechanisms or actions without developing explicit, reductive accounts of the sort of individuation at issue.

I have suggested a ‘holistic’ methodology whereby one relies on certain ‘clear’ intuitions about mechanism individuation in provisionally applying the account of guidance control, and then sees where this leads.⁹ As the theory gets applied to various cases, one will see how mechanisms must be individuated, if the account is to be plausible (that is, if the account is to match our considered judgments about moral responsibility). I do not think this is in itself disturbing or problematic. What *would* indisputably be problematic is if one individuated mechanisms in inconsistent ways—wiggling one way in one example and another way in another example, as Robert Nozick once put a similar point—simply to get the ‘right’ results (results that match our considered judgments about moral responsibility). In this instance the theory would not be generating the intuitions, but quite the reverse.

Judisch says:

Consider the case of Chum. Aside from the unfortunate name (what *were* his parents thinking?), Chum is a perfectly ordinary adult.¹⁰ He was raised in a happy home, received a normal moral education, and took responsibility (when a young man) for his mechanisms of ordinary practical reasoning and the like. Never once was he subjected to the slightest form of manipulation, and his moral and social development has left him a well-adjusted and responsible man. Now suppose that one night, while Chum is soundly asleep, he spontaneously develops a debilitating brain lesion. The lesion is situated in his neural network in such a way that his capacity for practical reasoning is severely impaired—the relevant mechanism no longer even approximates the standards of moderate reasons-responsiveness. Imagine now that a benevolent neuroscientist, Dr. White, is somehow made aware of Chum’s plight. Without rousing Chum, he springs into action. Unfortunately, there is no way he can remove the lesion without causing irreparable damage to Chum’s brain, but White has a few handy electronic devices that enable him (literally) to get around that problem. Here is how the devices work. The first, placed just ‘upstream’ of the lesion, takes as inputs the messages sent through the neural pathways headed right for the spot where the lesion is located, and it transmits the incoming

data via radio signals to the other device located just 'downstream' of the lesion, which device, in turn, relays the appropriate impulses to the neural pathways just downstream of it. The result is that the lesion is both successfully isolated and bypassed, its potentially deleterious effects completely cut off from the rest of Chum's brain; indeed, Chum's post-surgery cognitive architecture is *functionally equivalent* to his pre-surgery brain . . . When Chum awakens, he is of course completely unaware of the evening's events: as far as he is concerned, it is business as usual. (Judisch 2005)

I would seek to reply to Judisch's powerful critique in the holistic spirit adumbrated above. Judisch's view is that Chum can legitimately be considered morally responsible for his subsequent behavior. He contends however that on the Fischer/Ravizza approach, Chum cannot be deemed morally responsible. After all, we contend that if scientists manipulate your brain in a clandestine fashion, you have not taken responsibility for the mechanism that actually issues in your behavior (the 'manipulation-mechanism'), and it seems to Judisch that we should say the same thing about the case of Chum.¹¹

This is a difficult challenge. I am inclined to distinguish various kinds of intervention by neuroscientists. Neuroscientists can covertly implant desires (inputs); this gives rise to the problem of 'induced desires.' Alternatively, the neuroscientists can influence the ways inputs are processed and lead to outputs. In both of these cases I think the neuroscientists bring it about that a *different* mechanism operates—a mechanism different from 'ordinary practical reasoning.' Finally, neuroscientists may replace a biological part of the brain with another part, a part of the same or slightly different biological materials, or perhaps even made of synthetic or artificial materials. As I understand Judisch's thought-experiment, this latter sort of intervention is roughly what he has in mind. His intriguing example involves the analogue of a heart 'bypass' operation, but in the brain. An artificial mechanism is installed which allows the brain to function in the same way it has always functioned, but the electrical impulses bypass a problematic area. It functions in the same way both from the phenomenological perspective of Chum, and also in the sense that the 'function' from inputs to outputs is the same (even if it is realized in a slightly different way). Thus, I am inclined to think that this sort of intervention (in contrast to the first two) does *not* install a different kind of mechanism.

After all, if neuroscientists secretly installed a physically different part of the brain that functioned equivalently to the biological part, this would surely *not* create a different kind of mechanism. As far as I'm concerned, this would be like replacing one's carburetor with a functionally identical carburetor but made of slightly more durable material; it would still be the same kind of engine (for most conceivable purposes). Tinkering with the inputs, or with the way the inputs are processed (in the sense of modifying the function from the inputs to outputs) would indeed change the mechanism; but merely changing the physical realization of the processing that goes on (without changing the inputs or the processing of those inputs) does *not*, it seems to me, change the mechanism-kind.

Of course, I am employing the holistic strategy sketched above. It would perhaps be implausible to suppose that on some sort of independent grounds it is obvious that the first two sorts of intervention would install new or different mechanism-kinds, whereas the third sort of intervention would not. Rather, we can learn from a careful consideration of a range of examples, and also we can be guided by an attempt to secure a match with our considered judgments about moral responsibility. The question then becomes whether there is some intuitive plausibility to the emerging way of individuating mechanisms,

and, perhaps more importantly, whether this way of individuating mechanisms can be *generalized* to a broader range of cases in a *consistent* way.

The Transfer of Nonresponsibility

As Daniel Speak points out, my co-author and I have offered a critique of the modal principle (the Transfer of Nonresponsibility Principle) typically employed in the direct argument for incompatibilism about causal determinism and moral responsibility (Fischer and Ravizza 1994).¹² The direct argument does not proceed via the Principle of Alternative Possibilities, that is, the principle that moral responsibility requires the sort of control that involves genuine metaphysical access to alternative possibilities (regulative control). There are various versions of the direct argument, and one family of such arguments invokes a modal transfer of principle—a transfer of nonresponsibility principle.

Given plausible counterexamples to the initial version of the Transfer of Nonresponsibility Principle, one might try to modify it or restrict its application. Eleonore Stump has suggested that it ought to be modified or restricted so as to apply only to contexts in which there is just one causal pathway to the result in question; more specifically, these are cases in which there is no simultaneous or pre-emptive causal overdetermination. I have conceded that such a modification (or restriction) would evidently allow for the resulting principle to avoid the counterexamples. But I (and my co-author) have contended that problems remain. Speak says:

The . . . reason to be suspicious about Fischer and Ravizza's reply is that it seems to cast Stump's argument in an unfair light. Their claim is that Stump's restriction of the scope of Transfer NR to one-path cases is *ad hoc*. Presumably, this is because, as they insist in several places, the restriction is specifically designed to avoid counterexamples. But, again, this is surely an unfair characterization of Stump's motivation. It is true that the restriction serves to insulate the transfer principle from the sorts of examples constructed by Fischer and Ravizza. And it is also true that this is the result that Stump wants. She has not, however, insisted on the restriction merely to get this result. The restriction is motivated independently of her desire to generate a certain outcome. The motivation, of course, is to establish the appropriate context for the debate about determinism and moral responsibility. Stump's argument is that the contexts in which Transfer NR fails are irrelevant to the question at issue. The restriction on scope is an attempt to situate the relevant transfer principle in its proper context. Furthermore, this restriction is a plausible one, given Stump's argument. Fischer and Ravizza might have reasons for thinking that this restriction is inappropriate. But then these reasons should be brought forth. It will not do simply to insist that Stump has failed to meet her burden. (Speak 2005)

I agree that the motivation of the person who presents and then modifies a principle is one thing, and the plausibility of the principle quite another. All I suggested in previous work was that certain motivations may make it prudent to scrutinize the principle with some care. But the latter is really the important thing.

My main objection is that the principle, revised as suggested (so as to apply only to one-path cases), leads to bizarre and implausible results.¹³ Clearly, in a causally deterministic world, there will in fact (or may well be) contexts of simultaneous or pre-emptive causal overdetermination. Bizarrely, the One-Path Transfer Principle does not entail that the relevant agents are not morally responsible in those contexts. So it does not entail

incompatibilism. I would have thought that incompatibilism was the doctrine that causal determination rules out moral responsibility; but all we have thus far (given One-Path Transfer) is that causal determination in some (but not all) contexts rules out moral responsibility. One-Path Transfer does not yield incompatibilism!

Additionally, One-Path Transfer seems to give rise to a problem analogous to the idea that 'two wrongs make a right.' To see this, suppose that we have a one-path causally deterministic sequence. One-Path Transfer yields the result that the agent is not morally responsible in this context. But now add a second causally deterministic path to the same result. If the agent is not morally responsible in the first sequence, surely he must not be morally responsible when we add another causally deterministic path to the same upshot (given incompatibilist intuitions). But oddly Transfer One-Path does *not* yield the result that the agent is not morally responsible here. If anything, the situation seems *worse* for moral responsibility when there are *two* causally deterministic paths leading to the same upshot; certainly, it must seem no better from an incompatibilistic perspective. But One-Path Transfer fails to yield the result that the agent is not morally responsible in the two-path case. It is as if two metaphysical wrongs make a right!¹⁴

Conclusion

I wish again to express my heartfelt gratitude for the thoughtful, sympathetic, and probing critical discussions in this symposium. I have been fortunate to have such generous and gracious critics here (and throughout my career). I have sought to address some of the points that have been raised, but I have not been able to scratch the surface. Perhaps in some instances my preliminary attempts to reply to problems will have created other (and even worse) problems! In any case, my hope here (and always in my work) has been to get the view 'out there' in a clear and honest way. Even if it is ultimately rejected, I have tried to present a view that crystallizes a set of intuitive ideas about freedom, agency, and responsibility. I have sought to sketch an approach that can illuminate traditional debates about these central topics, and perhaps take the discussions in new directions.

NOTES

1. For a development of this contention, see Fischer and Ravizza (1994). See also Fischer (1994).
2. For some elaboration, see Fischer (1994, forthcoming).
3. McKenna goes on to point out that Ravizza and I have certain resources available on the basis of which to reply to this problem, and he makes some helpful and insightful suggestions of his own. (No, I will not criticize McKenna's sympathetic attempt to respond on my behalf!!)
4. See Van Inwagen (1983). For Lehrer's views, see, for example, Lehrer (1976, 1980).
5. For a detailed discussion of these matters, see Fischer (1979).
6. My point here is simply to distinguish the ability to do otherwise, in the 'all-in' sense of 'can,' from the possession of a general capacity (even one that cannot, in the circumstances, be exercised). I am not also contending that the pianist is in fact morally responsible for refraining from playing-this depends on whether he is aware of being chained to his desk, and the role these chains play in his deliberations.
7. I develop this point in Fischer (2002).

8. See, especially, Fischer and Ravizza (1998).
9. There is a more detailed development and defense of this sort of approach in Fischer (2004a).
10. Amusingly, my father-in-law's (now deceased) former wife was nicknamed, 'Chum!' She seemd to like this nickname.
11. I understand that the United States Central Intelligence Agency (CIA) makes a distinction between 'covert' and 'clandestine' operations. In a covert operation, we are allowed to know that some intervention is taking place, but the source of the intervention is secret. In a clandestine operation, even the fact of the intervention is not supposed to be known. For me, 'covert' and 'clandestine' interventions are not phenomenologically detectable by the relevant agent, and he is otherwise unaware of the intervention.
12. There are also discussions of this principle in Fischer and Ravizza (2000: part of a book symposium on J. M. Fischer and M. Ravizza, *Responsibility and Control: A Theory of Moral Responsibility*) and Fischer (2004b).
13. I have spelled these out in some detail in Fischer (2004b).
14. Of course, an incompatibilist may seek to supplement Transfer One-Path with other ingredients in order to generate the desired general result; but see Fischer (2004b).

REFERENCES

- FISCHER, JOHN MARTIN. 1979. Lehrer's new move: 'Can' in theory and practice. *Theoria* 45: 49–62.
- . 1994. *The metaphysics of free will: An essay on control*. Oxford: Blackwell Publishers.
- . 2002. Frankfurt-type compatibilism. In *Contours of agency: Essays on themes from Harry Frankfurt*, edited by S. Buss and L. Overton. Cambridge, Mass.: MIT Press.
- . 2004a. Responsibility and manipulation. *Journal of Ethics* 8: 145–77.
- . 2004b. The transfer of non-responsibility. In *Freedom and determinism*, edited by J. Campbell, M. O'Rourke and D. Shier. Vol. II of *Topics in contemporary philosophy series*. Cambridge, Mass.: MIT Press.
- . Forthcoming. A framework for moral responsibility. In *My way: Essays on moral responsibility*. New York: Oxford University Press.
- FISCHER, JOHN MARTIN, and MARK RAVIZZA. 1998. *Responsibility and control: A theory of moral responsibility*. Cambridge: Cambridge University Press.
- . 2000. Reply to Stump. *Philosophy and Phenomenological Research* 61: 477–80.
- GINET, CARL. 1990. *On action*. Cambridge: Cambridge University Press
- JUDISCH, NEAL. 2005. Responsibility, manipulation and ownership: Reflections on the Fischer/Ravizza program. *Philosophical Explorations* 8 (2): 000–000 (this issue).
- LEHRER, KEITH. 1976. 'Can' in theory and practice: A possible worlds analysis. In *Action theory*, edited by M. Brand and D. Walton. Dordrecht: Reidel.
- . 1980. Preferences, conditionals and freedom. In *Time and cause*, edited by P. Van Inwagen. Dordrecht: Reidel
- MCKENNA, MICHAEL. 2005. Reasons reactivity and incompatibilist intuitions. *Philosophical Explorations* 8 (2): 000–000 (this issue).
- SHABO, SETH. 2005. Fischer and Ravizza on history and ownership. *Philosophical Explorations* 8 (2): 000–000 (this issue).
- SPEAK, DANIEL. 2005. Semi-compatibilism and stalemate. *Philosophical Explorations* 8 (2): 000–000 (this issue).

156 JOHN MARTIN FISCHER

STRAWSON, GALEN. 1986. *Freedom and belief*. Oxford: Oxford University Press.

VAN INWAGEN, PETER. 1983. *An essay on free will*. Oxford: Clarendon Press.

WALLACE, R. JAY. 1994. *Responsibility and the moral sentiments*. Cambridge, Mass.: Harvard University Press.

John Martin Fischer, Department of Philosophy, University of California, Riverside,
California 92521, USA. E-mail: John.Fischer@ucr.edu