

THE FREE WILL REVOLUTION (CONTINUED)*

(Received 10 September 2005; accepted in revised form 5 October 2005)

ABSTRACT. I seek to reply to the thoughtful and penetrating comments by William Rowe, Alfred Mele, Carl Ginet, and Ishtiyaque Haji. In the process, I hope that my overall approach to free will and moral responsibility is thrown into clearer relief. I make some suggestions as to future directions of research in these areas.

KEY WORDS: free will, guidance control, Carl Ginet, Ishtiyaque Haji, Alfred Mele, moral responsibility, ownership of actual-sequence mechanism, overall framework for moral responsibility, Mark Ravizza, regulative control, reasons-responsiveness, semicompatibilism, William Rowe

I got some 'splainin to do ... (Ricky Ricardo [Desi Arnaz], *I Love Lucy*)

After having read the papers by William Rowe, Alfred Mele, Ishtiyaque Haji, and Carl Ginet, I am reminded of a story about Mark Twain. Twain was having a conversation with his wife, in which she (allegedly) told him (no doubt, “yet again”), “You will never amount to anything.” To which he replied, “Well, I can always be a bad example for others.”¹

Perhaps my views can at least be helpful insofar as they organize some intuitive ideas about free will and moral responsibility in a way that makes their structure – and even their problems – more perspicuous. I have had the goal of presenting an overall “framework” for moral responsibility. This framework involves (among other elements) an argument that moral responsibility does not presuppose free will in the sense of genuine access to

* I thank Michael McKenna for his thoughtful comments, his generous introduction, and all his hard work in putting this together. Also, I wish to thank J. Angelo Corlett for his very pleasantly surprising decision to do this volume, and for his outstanding work in editing this wonderful journal from its inception.

¹ This story was told by Mort Sahl in an interview with Terry Gross on her National Public Radio program, “Fresh Air.” The interview was conducted in December, 2003, and the program aired December 30, 2004, on WHYY in Philadelphia (<http://www.npr.org/templates/story/story.php?storyId=4251774>). Accessed on 28 July 2005.

metaphysically available alternative possibilities, a sketch of an account of the conditions of moral responsibility, and a suggestion about the “value” of acting in such a way as to be morally responsible. I have hoped to get this framework “out there” in a clear and forceful way as possible, with the thought that it will be helpful even to those who wish to disagree.

I am extremely grateful for the thoughtful, penetrating, and constructive papers by Rowe, Mele, Haji, and Ginet. I have learned a great deal from thinking about them, and I regret that I can only address a small subset of the genuine and pressing worries they each raise. This is in part due to limitations of space, but also in part due to limitations of my imagination. I can at least take some solace in the evident fact that I (or my views) have been a bad example for others!

In what follows I shall very briefly present an overview of the framework for moral responsibility I have offered in previous work, and then I shall select a few salient points to address from each paper. The principle of selection has more to do with what I find I can profitably say something about here than with the importance or force of the criticisms. I hope to have the opportunity (and ability) to address more of the criticisms in future work.²

1. SEMICOMPATIBILISM: AN OVERVIEW

I distinguish (in a pre-Quinean spirit) between the concept of moral responsibility and the conditions of its application. There are various possible specifications of the concept, including the “moral ledger” view and the Strawsonian view (according to which, roughly speaking, being morally responsible is being an apt candidate for what Peter Strawson called the “reactive attitudes”). I have not argued for any particular specification of the concept, although I

² For additional thoughts, see John Martin Fischer, “The Free Will Revolution,” *Philosophical Explorations* 8 (2005), pp. 145–155. The critical papers to which I was responding in this book symposium are: Daniel Speak, “Semi-Compatibilism and Stalemate,” *Philosophical Explorations* 8 (2005), pp. 95–102; Seth Shabo, “Fischer and Ravizza on History and Ownership,” *Philosophical Explorations* 8 (2005), pp. 103–114; Neal Judisch, “Fischer and Ravizza on History and Ownership: Reflections on the Fischer/Ravizza Program,” *Philosophical Explorations* 8 (2005), pp. 115–130; and Michael McKenna, “Reasons Reactivity and Incompatibilist Intuitions,” *Philosophical Explorations* 8 (2005), pp. 131–144.

have sometimes adopted a Strawsonian view as a “working hypothesis.” It may be that no particular specification uniquely and completely specifies our concept of moral responsibility; “moral responsibility” may thus be something like what Ludwig Wittgenstein called a “family-resemblance” term.

In contrast, I have defended a sketch of an account of the freedom-relevant (as opposed to epistemic) conditions of the application of the concept of moral responsibility. Crucial to my account is a distinction between two kinds of control: regulative and guidance control. An agent who has regulative control has genuine (whatever that is!) metaphysical access to alternative possibilities – he could have chosen and done otherwise. An agent who exhibits guidance control need not thereby possess regulative control. Such an agent possess a distinctive kind of “actual-sequence” control; he guides his behavior in a certain characteristic way.

A first step in my overall argument is to provide a plausibility-argument that moral responsibility does not require regulative control, but only guidance control. Perhaps it is better to characterize this step as offering a set of intuitive considerations, rather than, strictly speaking, an argument.³ I then go on to give an account of guidance control in terms of its two main ingredients: mechanism-ownership and moderate reasons-responsiveness. That is to say, an agent exhibits guidance control of an action insofar as it issues from his own, moderately reasons-responsive mechanism.⁴

On the view I have suggested, an agent becomes morally responsible by “taking responsibility” – he acquires control (in part) by taking control. This process of taking responsibility involves acquiring certain beliefs. More specifically, in the gradual process of moral education an agent would typically acquire the beliefs that he can differentially affect his environment based on his choices and bodily movements, and that he is a fair target for certain morally-charged reactions based on his choices, bodily movements, and their effects on the environment. When an individual completes the process of acquiring these beliefs, he makes certain mechanisms of action “his own.” Given that certain

³ See, for example, John Martin Fischer, “Responsibility and Control,” *The Journal of Philosophy* 79 (1982), pp. 24–40; and John Martin Fischer, *The Metaphysics of Free Will* (Oxford: Blackwell Publishers, 1994).

⁴ John Martin Fischer and Mark Ravizza, *Responsibility and Control* (Cambridge: Cambridge University Press, 1998).

additional conditions are satisfied, such an agent can truly say, “I did it my way.”⁵

In order to have guidance control of behavior, an agent’s mechanism must be appropriately responsive to reasons (that is, “moderately responsive to reasons”). Moderate reasons-responsiveness involves a distinction in the nature and structure of an agent’s recognition of reasons and his reactions to them; more structure (or perhaps a different “profile”) is required in the reasons-recognition component than the reasons-reactivity component of moderate reasons-responsiveness.⁶ Moderate reasons-responsiveness is defined in terms of modal or dispositional properties, and thus facts about non-actual possible worlds are relevant. But it is a distinctive feature of my account that these facts are relevant *not* in virtue of indicating genuine access to other possible worlds, but simply in virtue of helping to specify modal or dispositional properties *of the actual sequence of events*.

Since moral responsibility does not require “could have done otherwise” (regulative control), the semicompatibilist can sidestep traditionally vexatious problems about the relationship between (say) God’s omniscience or causal determinism and freedom to do otherwise. The distinctive stance – “revolution” is no doubt unduly pretentious – of semicompatibilism is to prescind from questions about regulative control, and to focus on the “actual sequence”

⁵ For further explanation and development of this point, see John Martin Fischer, *My Way* (Oxford: Oxford University Press, 2006). In her wonderful short piece, “Ixnay on the My Way,” written in 1997, Sarah Vowell writes:

Is there anything nicer than a really good TV obituary? Any day now, Peter Jennings will cut away from some freak mudslide story (casualties: six registered voters), face another camera, and announce Frank Sinatra’s death. Later, the *World News Tonight* credits will roll over a tasteful montage of Frank’s film stills and album covers. The other networks will run similar tributes, as will the brainiacs at *Entertainment Tonight* and those swingers on *The NewsHour* at PBS. But you know what? It will not matter whether Sinatra’s video wake is hosted by the tweedy Jim Lehrer or the perky Katie Couric. Because each and every remembrance will be accompanied by the same damn song: the most obvious, unsubtle, disconcertingly dictatorial chestnut in the old man’s vast and dazzling backlog [Sarah Vowell, “My Way,” in Sarah Vowell, *Take the Cannoli* (New York: Simon and Schuster, 2000), p. 160]. There are interesting discussions of Kitty Kelly’s notorious *His Way* and a fanzine, *Our Way: In Honor of Frank Sinatra*, in Vowell, *Take the Cannoli*, pp. 71–80.

⁶ Fischer and Ravizza, *Responsibility and Control*.

leading to the behavior under evaluation. More specifically, it argues that (say) causal determinism is compatible with moral responsibility, quite apart from whether causal determinism is compatible with regulative control. An important presupposition of this argument is that the case for the incompatibility of causal determinism and regulative control is *different from* and *considerably stronger than* the case for the incompatibility of causal determination in the actual sequence and moral responsibility.

I (and my co-author) have attempted to give a unified, systematic account that ties together guidance control of actions, omissions, and consequences (particular and “universal”), as well as character traits and emotions. The systematic and unified nature of the account – employing the same fundamental ingredients of control and natural extrapolations of them – provides reason to accept the basic association of moral responsibility with control, as well as the particular accounts of guidance control of actions, omissions, consequences, and emotions.

Two salient features of our account of guidance control are that it is “subjective” and “historical.” On the account, an agent can possess guidance control and thus be morally responsible only if he sees himself in a certain way and thus has taken responsibility for the relevant mechanism (from which the behavior in question issues).

Finally, I distinguish different accounts of the “value” of acting freely and thus being morally responsible. Clearly, we assign some value to acting freely (even if this is not a hegemonic value). That is, we prefer to be the sorts of creatures who at least sometimes act freely, rather than robots or mere automata or even sentient creatures who lack free will (such as nonhuman animals). On the approach that prizes regulative control, the value of exhibiting the relevant sort of control (in virtue of which we can legitimately be held morally responsible) is that we *make a difference to the world*; on this approach, the value of being morally responsible is the value of making a difference.

In contrast, I contend that the value of moral responsibility is the value of *making a certain sort of statement* – the value of creative self-expression. If what we value in being morally responsible is a kind of artistic self-expression – writing a sentence in the narrative of our lives – then this provides further reason to suppose that moral responsibility does not require regulative

control, but merely guidance control. This is because the relevant kind of artistic self-expression need not involve access to alternative possibilities.⁷

2. REPLY TO ROWE

An important part of the framework I have sketched is a plausibility-argument (or set of intuitive considerations) for the idea that moral responsibility is a matter solely of how the actual sequence of events unfolds – genuine access to appropriate alternative possibilities is not necessary. I thus reject the Garden of Forking Paths (regulative control) model of moral responsibility.

It has admittedly been difficult to articulate the argument in a satisfactory way; perhaps this is why I have tried to improve it (or at least “fill it in”) on several occasions! In order better to understand the structure of my view here, it will be helpful to lay out the nub of Rowe’s worry. He points out that even in a Frankfurt-type case such as Rowe’s Case 3, the agent has a certain sort of alternative possibility: the power not to agent-cause the volition to act. He recognizes that I would respond that this sort of alternative is insufficiently robust to ground moral responsibility attributions: it lacks sufficient Oomph. My point, of course, does not involve accepting a regulative-control model of moral responsibility; rather, the point is that if one accepts such a model, the alternative possibilities in question will not do the trick. This is a point with which I am in agreement with the libertarian, Robert Kane, who insists on the requirement of “plural control” for moral responsibility.⁸

Rowe characterizes my position as follows:

But when we look very carefully at what Fischer says, it becomes clear that he so uses the expression “Oomph” that unless the alternative includes a *volition* or an *intention* to do something else (or to do nothing) then it follows *by definition* that there is no alternative present with sufficient “Oomph” to ground moral responsibility. And this being so, Fischer wins the argument. ... In response to Reid’s view, what Fischer has done is to simply *legislate* that unless the alternative involves a *volition* not to act, the alternative is not “robust enough” to render that agent

⁷ John Martin Fischer, “Responsibility and Self-Expression,” *The Journal of Ethics* 3 (1999), pp. 277–297.

⁸ Robert Kane, *The Significance of Free Will* (Oxford: Oxford University Press, 1996), pp. 133–135.

morally responsible. Given this nifty piece of legislation on Fischer's part, it comes as no surprise that he is able to conclude that Reid's agent cause account of a volition act is not "sufficiently robust" and, therefore, lacks the "Oomph" required to render the agent morally responsible for not agent causing a volition to act. So, of course, Reid's theory of agent causation does not have enough "Oomph" to satisfy Fischer. For Fischer has so defined "Oomph" that it is logically impossible for Reid's agent causal account of a free act to satisfy Fischer's definition. What Fischer has not shown is that Reid's theory of freedom is inadequate to ground ascriptions of moral responsibility. All that he has done is *define* "Oomph" so as to preclude a person's agent causing a volition to act from possessing Oomph. There is nothing right or wrong in using the expression "Oomph" in that way. But what Fischer needs to do is to provide a serious *argument* for the conclusion that when a person agent causes a volition to perform some act, and does perform that act, it is not an act for which the agent can be morally responsible. When, and if, he succeeds in doing that he will have provided some rational support for his rejection of this, and other incompatibilist views of moral responsibility.⁹

I would contend that more has been said, and that less needs to be said, than Rowe supposes. But first I would point out that Rowe is in good company in supposing that I have simply sought to issue an *ex cathedra* pronouncement on these matters. R. Jay Wallace (a compatibilist) writes:

[Fischer's] response seems suspiciously ad hoc, however. After all, there is surely *some* sense in which the agents in Frankfurt's scenarios cannot do otherwise; it almost looks as if Fischer has singled out that sense and simply declared it to be the sense that matters to the debate. At the least, we need a convincing and principled reason [to adopt Fischer's view of the sort of alternative possibilities required for moral responsibility].¹⁰

Well, I suppose I would rather be "nifty" than "suspiciously ad hoc"! I said above that more has been said about these matters, and that less needs to be said, than Rowe supposes. The same is clearly true of Wallace. To start with the "less needs to be said" part, I would point to what must be a truism in philosophy, namely, that not everything (or even, lamentably, everything important) can be *argued for*. One sometimes has to rely on "considered judgments" or "intuitions" or basic, fundamental insights that simply cannot be further explained. So I reject the idea that I "need" to provide a "serious *argument*" for my view here, or that "at the least, we need a convincing and principled reason [to adopt my view]." Does Kane offer a "serious

⁹ William Rowe, "Free Will, Moral Responsibility, and the Problem of 'Oomph'," this issue of *The Journal of Ethics*.

¹⁰ R. Jay Wallace, *Responsibility and the Moral Sentiments* (Cambridge: Harvard University Press, 1994), p. 262.

argument” for his plural control requirement on moral responsibility? How many times in philosophy does one actually get a “convincing and principled reason” for a highly disputed and contentious proposition? Is it not more typical that one gets at most plausibility arguments and suggestive intuitive considerations? I would have thought that it would be more accurate and reasonable to say that *the very most* we could ever expect in this dialectical terrain would be a convincing and principled reason to accept a highly contentious proposition!

Now to the “more has been said” part. Although I concede that my arguments have been preliminary, partial, and somewhat tentative, I am puzzled as to why both Rowe and Wallace seem to suppose that I have not offered any sort of argument for the view that *robust* alternatives – alternatives with sufficient Oomph (properly construed) – would be required to ground moral responsibility, on the regulative control model. Here I shall simply adumbrate some of the considerations I have sought to develop in previous work.

I begin with the distinction, well known to anyone who has considered the worries about libertarian agency, between mere possibility and ability. That is, there is an important distinction between the mere possibility that something different should occur (perhaps by accident or as a result of some random process) and an agent’s having the ability to do otherwise (or bring about a different outcome). Ability is not the same as mere possibility, and this same problem (that afflicts certain libertarian accounts of agency) also faces anyone who supposes that the mere possibility of something different’s occurring is sufficient to ground moral responsibility attributions. How can such a possibility be sufficiently substantive to support the idea that an agent *controls* his choices and behavior, and is morally accountable for them? I have used the metaphor of “alchemy” here; seeking to get moral responsibility out of such exiguous alternatives seems to be like trying to get gold from straw.

I have also used what might be called a “No-Difference” kind of argument. This sort of argument presents a certain scenario and elicits the reader’s view of or judgment about the scenario. Then it *changes* the scenario by adding (or subtracting) some crucial element, and it claims that the change *makes no difference*. The difference in question thus is alleged not to make a difference; the change is no help to the proponent of the view under consideration.

So I have presented the following No-Difference argument. Imagine a context in which an agent has no alternative possibilities

(perhaps in virtue of the truth of causal determinism, supposing that causal determinism is incompatible with the relevant sort of alternative possibilities as per Van Inwagen's Consequence Argument or some similar argument). *Ad arguendo*, suppose, further, that in virtue of lacking alternative possibilities the agent is not morally responsible. Now add to the scenario an alternative sequence in which there is no voluntariness at all – no Oomph. That is, add a scenario in which the agent merely involuntarily blushes or furrows his brow or exhibits some esoteric (to the agent) neurological pattern, which triggers a substantial intervention by a “nefarious” or even “nice” neurosurgeon. I have suggested that a thoughtful and fair-minded and reasonable person would say that merely adding this sort of alternative possibility could not possibly make a difference as to the agent's moral responsibility. More carefully, I have suggested that merely adding this sort of alternative possibility could not *in itself and apart from indicating something about the actual-sequence* make a difference to the agent's moral responsibility.¹¹ To suppose that the availability of such a “wimpy” alternative possibility could transform a context of no moral responsibility into one of moral responsibility is tantamount to a belief in alchemy, I suggested. More carefully, it just seems very implausible to suppose that the existence of such alternative possibilities *in themselves and apart from pointing to causal gaps in the actual sequence* could make the difference in question (Whether causal gaps in the actual sequence are required for moral responsibility is quite another question – one which has not escaped my attention).¹²

My No-Difference style of argumentation here is structurally similar to a strategy of argumentation employed extensively (and to good effect) by Randolph Clarke.¹³ To oversimplify, Clarke begins by imagining that an agent in a causally deterministic world exercises what he calls “direct active control.” This involves acting as a result of reasons (in an appropriate--nondeviant--way). Of course, this control must not be thought to require alternative possibilities or “sourcehood” defined indeterministically. Now Clarke imagines that one subtracts causal determination and inserts mere indeterministic

¹¹ For a more careful treatment of this point, see John Martin Fischer, “Frankfurt-type Compatibilism,” in S. Buss and L. Overton (eds.), *Contours of Agency: Essays on Themes from Harry Frankfurt* (Cambridge: The MIT Press, 2002), pp. 1–26.

¹² Fischer, *The Metaphysics of Free Will*.

¹³ Randolph Clarke, *Libertarian Accounts of Free Will* (Oxford: Oxford University Press, 2003), pp. 74–82.

event-causation of a certain sort. He contends that the distinctive features characteristic of the relevant kind of control (“direct, active control”) would still be present, and thus the envisaged difference would not make a difference. This sort of No-Difference strategy of argumentation is structurally similar to the argument I have offered for my contention that it is guidance-control, and not regulative control, that is the freedom-relevant condition for moral responsibility.

I do not suppose that I have offered a decisive argument against the requirement of regulative control for moral responsibility. But I should have thought that such an argument would not be available, and I have done my best to present some intuitive considerations that might move a reasonable and fair-minded person to reject the requirement.

In the end, however, I am not satisfied that I have fully understood or done justice to Rowe’s objection to my views. I want to try once more. Perhaps Rowe is willing to concede that I have in fact offered arguments (of the inconclusive but suggestive sort mentioned above) for the claim that alternative possibilities must be robust – must have Oomph – in order to help to ground moral responsibility attributions, on the regulative control model. His point might instead be that I have somehow simply “stipulated” that “Oomph” (in the relevant sense) must involve voluntariness, and in certain contexts such as Case 3 (or perhaps even all contexts), Reid would have it that the alternative sequence does not involve a volition (and thus voluntariness). For Rowe (as opposed perhaps to Wallace), the problem is not that I have not plumped for Oomph; the problem is that I have simply legislated that Oomph must involve voluntariness, and this seems inconsistent with Reid’s account of how agent causation works – an account Rowe finds at least attractive and worth taking seriously.

I am still not sure I see the exact form of the objection. There certainly may be contexts in which there are good reasons to do various things, and, in such a context, it would seem that Reid or in general an agent-causationist would say that the relevant agent has the power to agent-cause the volition to *X*, while also having the power to agent-cause a volition to do some incompatible thing *Y*. So I do not see why it would be a matter of definition for Reid that the alternative scenario (or range of such scenarios) would lack voluntariness. Now in a Frankfurt-type case such as Case 3, the alternative scenario does lack voluntariness and indeed any kind of intentional

behavior by the agent. Upon careful reflection, it is my intuition that the existence of this sequence does not in itself and apart from indicating something about the actual sequence confer moral responsibility on the agent (or even help to ground such responsibility). You can say that the reason is that the alternative sequence lacks voluntariness, or that it lacks intentional behavior, or whatever – the specific characterization is not important to me. It just seems that, however characterized, the existence of the alternative scenario is in itself too flimsy to help to explain or ground moral responsibility.

Let us be a bit more specific about this. Suppose that before he were to do anything but keep the money, Jones were to furrow his brow or evince an esoteric neurological pattern in his brain or perhaps blush red on his forehead. Seeing this, the Devil can intervene and ensure that Jones keep the money. How could it be in virtue of a mere unintentional “twitch” or “sign” such as the above that the agent is morally responsible? More carefully, it seems to me that the *only* way that the existence of alternative scenarios of this sort could help to explain or ground moral responsibility would be in virtue of pointing to something (indicating something) about the *actual sequence*. Of course, on Reid’s view, such alternative possibilities may indicate that Jones agent-caused his volition to keep the money – and *some* alternative possibilities are necessary for there to be such causation. So be it; but then the important issue is whether moral responsibility requires that the actual sequence involve agent-causation, *not* whether the existence of alternative possibilities (the presence of regulative control) is what *explains* and *grounds* (or even helps to explain or ground) moral responsibility.

3. REPLY TO MELE

Alfred Mele presents Phil, who has come under the baleful but persuasive influence of Ted Honderich in his trip to the lovely city of London (Phil evidently avoided other terrorists!). Phil meets all intuitive criteria for being a fully and robustly morally responsible agent, but he fails to meet one of the three criteria I (and Mark Ravizza) have set out for “taking responsibility,” and thus he fails to meet the “subjective” criteria for moral responsibility (on our approach). More specifically, since Phil does not see himself as an apt target of the reactive attitudes, he does not count as morally

responsible, on our approach. In contrast, the agoraphobe Fred *does* count as morally responsible, on our approach, despite the fact that Fred's agoraphobia is "so powerful that he has not ventured out of his house in ten years..." He counts as morally responsible, on our view, because (roughly speaking) he would leave his house (holding the relevant things fixed) if there were a raging fire. His psychological compulsion (issuing from the agoraphobia) is strong but not literally irresistible. Mele thinks we have got it backwards here, and that we should re-evaluate our views of the moral responsibility of such agents as Phil and Fred.

I agree with Mele that such agents as Phil and Fred present good challenges to the view of moral responsibility presented (and subsequently defended) by Ravizza and me. My general methodological disposition is to seek to capture the *clear* cases by appealing and intuitively natural principles, but to admit that these principles may well have jarring consequences in certain cases. Of course, on the sort of methodology I favor, there must be some sort of "reflective equilibrium" (to use John Rawls' phrase) in which one adjusts one's principles and "considered judgments" to seek harmony. Given the distinction between moral responsibility and blameworthiness/praiseworthiness, I frankly do not think it is *evident* that the Fischer/Ravizza approach to such agents as Phil and Fred counts decisively against our general theory. After all, the phenomena of moral responsibility are themselves messy around the edges, and it would be unreasonable to suppose that a largely successful and plausible approach would yield entirely comfortable results along all its perimeters.

But in the constructive spirit in which Mele's reflections are offered, I would also argue that I believe that I could accept Mele's views without in any way jeopardizing the main features of the Fischer/Ravizza approach: that moral responsibility is compatible with causal determinism, and even that moral responsibility is historical and subjective in the relevant sense, and a matter of the appropriate reasons-responsiveness of the agent's own behavior-producing mechanism.

To consider Phil first. The various individuals discussed by Mele – the addict who believes he is an addict and thus cannot in the end successfully resist the urge for the drug, the sailor who knows that his rudder is broken, and Phil – all lack a kind of "self-engagement." But Mele is right to want a finer-grained articulation of the self-engagement in question, and he is correct to

note that Phil's situation is distinctive. Here are the three conditions Ravizza and I have proposed for an agent's "taking responsibility" for the kind of mechanism issuing in the relevant behavior:

First, an individual must see himself as the source of his behavior ... in the sense that he must see that his choices and actions are efficacious in the world. ... Second, 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. ... The third condition on taking responsibility requires that the individual's view of himself specified in the first two conditions be based, in an appropriate way, on the evidence.¹⁴

These three conditions, suitably qualified and refined, are supposed to define a kind of "self-engagement" characteristic of a morally responsible agent. On the Fischer/Ravizza approach, they help to define "taking responsibility" and thus mechanism-ownership, one of the two chief elements of guidance control. It is notable that whereas the addict who knows he is an addict and the sailor who knows his rudder is broken (and other, similar individuals) fail to meet the first condition, Phil fails to meet the second condition. Thus, it may be that there is an important difference between Phil and the other agents, in virtue of which Phil is morally responsible whereas the others are not.

As I said above, I am not confident about what to say about this case, but I am willing to entertain dropping the second condition. On this approach, one would keep the first and third conditions as defining some subjective notion of self-engagement related to the intuitive idea of "taking responsibility"; without the second condition, this notion could not plausibly be thought to capture the commonsense notion of taking responsibility, but this should not pose a problem, as long as we are clear that the characterization does not purport to analyze the pre-theoretic notion of "taking responsibility." On this approach, an individual "takes responsibility" (in the special, theoretical sense) insofar as he sees himself as an agent in a distinctive way, that is, sees that his choices and actions are efficacious in the world – and this conception of himself is based on the evidence in the appropriate way. It should be clear that this

¹⁴ Fischer and Ravizza, *Responsibility and Control*, pp. 210, 211, and 213.

emendation preserves all the basic features of the Fischer/Ravizza approach, while allowing us to accommodate Mele's intuition about Phil.¹⁵

Some philosophers (and I am not including Mele in this group) have dismissed the Fischer/Ravizza Semicompatibilism out of hand because of its implications for individuals such as Phil. I hope that the above discussion shows that all of the major components of our theory can be maintained compatibly with an adjustment that shows that we need *not* be committed to these implications.

Similarly, I would be open to an adjustment of the sort Mele suggests with respect to Fred. That is, I believe that Ravizza and I could accept Mele's generous suggestions that "an attractive strategy for avoiding the (apparent) problem that I have been developing is to beef up the reasons-reactivity condition in such a way that Fred and agents with equally severe psychological maladies of the pertinent kind do not count as reasons-responsive enough to be morally responsible for the relevant behavior."¹⁶ This posits a more refined notion of moderate reasons-responsiveness, with what might be called "spheres of responsiveness;" the "outer spheres" would not necessarily indicate sufficient responsiveness for moral responsibility.¹⁷ Of course, it may not be straightforward to characterize precisely the "borders" of the spheres; that is, it might not be easy to say exactly what degree of strength of the relevant sort of urge renders the agent in question immune to moral responsibility. But this need not be my task here.

Ultimately, I am not sure what to say about the difficult cases Mele presents. My main reply is to point out that these *are* difficult cases, and, further, that the major components of my overall theory can be maintained, even with adjustments to accommodate different views about these cases. That is, one can accommodate Mele's views while continuing to maintain a guidance-control based,

¹⁵ Similarly, this adjustment in the theory would successfully respond to the thoughtful criticisms developed in Andrew Eshleman, "Being is Not Believing," *Australasian Journal of Philosophy* 79 (2001), pp. 479–490. I have toyed with this adjustment for awhile, and Carl Ginet also independently suggests it in his contribution to this issue of *The Journal of Ethics*.

¹⁶ Mele, "Fischer and Ravizza on Moral Responsibility," this issue of *The Journal of Ethics*.

¹⁷ Note that this is consistent with the not entirely uncontentious Fischer/Ravizza claim that reactivity is all of a piece; an agent who can react to any reason may have great difficulty in doing so in any particular context.

compatibilistic theory, according to which moral responsibility is an essentially historical (and even suitably subjective) phenomenon.

4. REPLY TO HAJI

Haji's critical piece is subtle and raises many fascinating issues not all of which I can tackle here; with apologies, I shall select a few salient points to discuss. One might profitably distinguish "forward-looking" from "backward-looking" aspects of agency. Backward-looking aspects include attributions of moral responsibility, whereas forward-looking aspects include deliberation, practical reasoning, and so forth. I have argued that genuine metaphysical access to alternative possibilities – regulative control – is required for neither aspect of agency.¹⁸ Haji distinguishes moral responsibility from a distinctive circle of "judgments of deontic morality" (including moral ought-judgments); he is willing to concede that moral responsibility, but not the judgments of deontic morality, are compatible with the lack of regulative control. Further, he finds fault with my views about the relationship between ought-judgments and their signature "action-guiding" function.

On my view, forward-looking dimensions of agency require epistemic openness of a certain sort, but not genuine metaphysical openness (or even the belief in such openness).¹⁹ On my view, then, ought-judgments can play their distinctive action-guiding role in the space of epistemically open options.

Haji objects, saying:

... this way of preserving the action-guiding role of morally deontic judgments incurs a cost. Suppose our world is causally determined and thus, no one can ever do other than what he or she in fact does. ... it appears that there will be many occasions on which people ought morally to do things that they cannot do. On Fischer's account of moral guidance, on all such occasions moral judgments will provide moral guidance only if the relevant agents are irrational in that they have false beliefs concerning what they can in fact do; on all such occasions, people will

¹⁸ For an extended discussion of the forward-looking facets of agency, together with an analysis of the literature surrounding the "Frankfurt-type examples," see John Martin Fischer, "Free Will and Moral Responsibility," in D. Copp (ed.), *Oxford Handbook on Ethical Theory* (Oxford: Oxford University Press, 2006), pp. 321–354.

¹⁹ Fischer, "Free Will and Moral Responsibility."

have to believe falsely that they can do what morality requires of them if morality is to guide them in the way in which Fischer proposes.²⁰

I would point out that it is perhaps a stretch to use the term “irrationality” for the epistemic situation of an agent in a causally deterministic world. Such an agent might even accept the truth of causal determinism and the conclusion of the Consequence Argument (that causal determinism is incompatible with regulative control), and thus he might know that, whatever it is that he chooses and does, that is the *only thing* he “can” (in the relevant sense) choose to do (and do). It does not follow that in advance he knows specifically what he will choose and do, and it thus does not follow that ought-judgments cannot provide guidance among the options that are “epistemically open” to the agent-open, for all he knows.

Now this picture of the role of ought-judgments does presuppose that the agent does not know in advance what specifically he will choose and do (in the relevant contexts). This is a lack of knowledge, but is it a form of “irrationality”? It is not clear to me that the agent must have *any* false belief; it seems to me that he must simply not believe certain things which are in fact true. Further, to suppose that this lack of knowledge counts as “irrationality” in an objectionable sense is highly contentious; after all, if the agent did in fact know in advance exactly what he would choose and do, then it is unclear that he could be genuinely “active” in his life at all. To suppose that it is problematic that an agent must lack specific knowledge of his future choices and behavior is highly implausible; precisely this sort of incompleteness in one’s knowledge is often thought to be a necessary condition of agency itself. It is thus completely unclear that the sort of epistemic openness I have invoked constitutes genuine “irrationality,” or any sort of uncontroversial epistemic defect in an agent.

Haji points out that I (and my co-author) have adopted a certain sort of holistic methodology, quoting us as follows:

We offer what we take to be strong plausibility-arguments for the claims that moral responsibility does not require alternative possibilities, and that causal determinism in itself does not rule out moral responsibility. We then offer a general theory of moral responsibility that shows how it is *possible* to defend, in detail, these views – in particular, that moral responsibility is compatible with causal determinism. This

²⁰ I. Haji, “Frankfurt-Type Examples, Obligation, and Responsibility,” this issue of *The Journal of Ethics*.

theory gains some credibility from its *systematic and unified* treatment of moral responsibility for actions, omissions, consequences, and even traits of character. Of course, our arguments for the overall approach are not *decisive*, and various elements remain to some degree or another vague and undeveloped.²¹

But, according to Haji, we are hoist by our own petard:

The strong and promising suggestion of Fischer's is that support for a controversial view or principle, one for which no direct arguments are decisive, can be marshaled by examining virtues of the *overall package*, in Fischer's instance, the overall actual sequence approach to responsibility.

Analogously, when assessing *K* [the maxim that "ought" implies "can"], I suggest that we look to overall accounts of the concept of moral obligation and then see what these accounts imply about the truth of *K*. Although this is not the place to delve into pertinent details, I have proposed that our most promising views of the concept of moral obligation – roughly, those Michael Zimmerman and Fred Feldman defend – include *K* as a theorem.²²

So Haji offers a "So's YOUR Momma" or, in Latin, *tu quoque* argument. That is, he contends that my preferred holistic methodology should issue in an acceptance of *K*. I reply that I would like to consider more carefully whether excellent overall theories of obligation, such as those of Michael Zimmerman and Fred Feldman, are *essentially* committed to some sort of "ought-implies can" maxim. It is most likely that Zimmerman and Feldman are not concerned with the implications of Frankfurt-type examples, or any other abstract metaphysical considerations having to do with the relationship between moral responsibility and regulative control, in developing their accounts of moral obligation (This is not to say that these theorists are not highly interested in such metaphysical issues, or have not addressed them in their work). It may be that the theories they develop could be *adjusted* so as to preserve the important core of theorems about obligation without also implying *K*. That is, once considerations about the relationship between moral responsibility and regulative control are introduced, it might be possible to adjust the theories so as to accommodate the core of the theories without also being committed to an "ought-implies-can" principle. It would only be if the most plausible theories of obligation could not be adjusted in this sort of way that the Overall Package methodology could be invoked against my views here.

²¹ Fischer and Ravizza, *Responsibility and Control*, as quoted in Haji, "Frankfurt-Type Examples, Obligation, and Responsibility," this issue of *The Journal of Ethics*.

²² Haji, "Frankfurt-Type Examples, Obligation, and Responsibility," this issue of *The Journal of Ethics*.

I have argued that, quite apart from issues pertaining to the range of deontic judgments, causal determinism would seem to be compatible with other significant moral assessments. More specifically, I have contended (following a suggestion of Clarke in conversation) that causal determinism would appear to be entirely compatible with one's having a sufficient reason to behave in a certain way.²³

Very graciously, Haji calls this a "striking proposal," but he also states, "I take exception to the claim, (F1) Individuals are morally blameworthy when they fail to do what they have sufficient reason to do, (assuming it is epistemically open for the agent to do what he has failed to do)."²⁴ Haji's subsequent discussion insightfully sorts through some of the intricacies in this dialectical neighborhood.

I am grateful to Haji for making it clear that what I should have said is that causal determinism seems to me to be entirely consistent with one's having a sufficient moral reason to behave in a certain way. I am inclined to think that one is indeed morally blameworthy insofar as one fails to do what one has sufficient moral reason to do, assuming that it is epistemically open for the agent to do what he has failed to do, and, perhaps, given that it was reasonable to expect him to recognize this sufficient moral reason. So, even if one banished all judgments of deontic morality from a causally deterministic world, one could still have moral responsibility, practical reasoning, normative guidance by ought-judgments, and blameworthiness stemming from failing to do what one has a sufficient moral reason to do. One might in such circumstances find that one does not miss the circle of deontic judgments much at all.²⁵

²³ John Martin Fischer, "'Ought-Implies-Can', Causal Determinism, and Moral Responsibility," *Analysis* 63 (2003), pp. 244–250; and Haji, "Frankfurt-Type Examples, Obligation, and Responsibility," this issue of *The Journal of Ethics*.

²⁴ Haji, "Frankfurt-Type Examples, Obligation, and Responsibility," this issue of *The Journal of Ethics*.

²⁵ It still is a mystery to me, as I indicated in Fischer, "'Ought' Implies 'Can,' Causal determinism, and Moral responsibility," what the relationship is between claims such as "S has a sufficient moral reason to X" and "S ought morally to X." It is sometimes thought that the former provides some sort of analysis of the latter; if so, then I am wrong to suppose that there is a difference in the entailments with respect to K. If, on the other hand, I am correct about the difference in the entailments, then the analysis is faulty.

5. REPLY TO GINET

Although all of the critical papers in this symposium are both penetrating and (excessively) kind, I am especially grateful to Ginet. Ginet was my dissertation supervisor at Cornell University, and he has continued to be a model of philosophical mentorship throughout my entire career. His paper beautifully displays his characteristic combination of critical acumen and constructive spirit – to which I am now, and have always been, deeply indebted.

Having developed a nice problem for the Fischer/Ravizza approach, Ginet writes:

I think there is an easy way for Fischer and Ravizza to avoid both this particular problem and general puzzlement about how to individuate action-producing mechanisms. Their specification of moderate reasons-responsiveness entails that the mechanism that produced the action be such as to make it the case that the *agent* would recognize and act on each of a suitable pattern of reasons to do otherwise were the agent to have that reason, i.e., that the mechanism *not* be such that the agent would *not* have thus responded. Putting it this way suggests an obvious solution, namely, instead of talking about a *mechanism's* being reason-responsive, to talk about the *agent's* being reasons-responsive at the time of the action, to require that the *agent* then was such that they would recognize and act on each of a suitable pattern of reasons to do otherwise were they to have that reason.²⁶

Ginet thinks that we can avoid the problems about mechanism-individuation by simply eliminating talk of mechanisms, and focusing on “the agent as he actually was:”

Given the way in which it is true of Frankfurt-type examples that the agent could not have done otherwise – there is present something that would intervene to make the agent do (or not do) the thing should he show any “sign” of not doing it (or of doing it) – it does not, as far as I can see, follow that the *agent* is not reasons-responsive. Everyone agrees that it's not the way the action (or non-action) actually came about that makes the agent not responsible. It seems that the *agent as he actually was*, the actual sequence agent – as distinguished from the way he would have been if the backup thing had intervened – *was* reasons-responsive. Retaining all the dispositional (i.e., counterfactual) properties he actually had, and not shifting to those he would have had in the alternative, intervention scenario, the agent would have recognized and reacted to sufficient reasons to do otherwise (of a suitable pattern) – i.e., he was moderately reasons-responsive.²⁷

²⁶ Ginet, “Working With Fischer and Ravizza’s Account of Moral Responsibility,” this issue of *The Journal of Ethics*.

²⁷ Ginet, “Working With Fischer and Ravizza’s Account of Moral Responsibility,” this issue of *The Journal of Ethics*.

This is a very nice suggestion. The Frankfurt-type examples are structurally similar to a range of other scenarios that describe objects whose dispositional or “counterfactual” properties would change in unusual or weird ways, due to (perhaps) bizarre external or internal factors. I have called such situations, “Schizophrenic Situations”, and have contended that they mark out a kind of “swerve in logical space.”²⁸ Frankfurt-type examples pertain to the “active” power of freedom, whereas other Schizophrenic Situations pertain to “passive powers,” such as solubility, and so forth.²⁹

These Schizophrenic Situations present many philosophical difficulties and puzzles. In analyzing the relevant power, the Fischer/Ravizza strategy “reaches into the actual sequence of events” and seeks to latch onto the actually-operative mechanism; having fixed on this mechanism, we seek to identify its counterfactual properties. Of course, this strategy faces various challenges, including the puzzles associated with mechanism-individuation. In contrast, Ginet suggests that we ought to seek to latch onto the “agent as he actually was,” and then seek to identify his counterfactual properties. Just as the counterfactual properties of the actual behavior-producing mechanism are different from those of the behavior-producing mechanism in a range of alternative scenarios, so the counterfactual properties of *the agent as he actually was* will be different from those of the agent as he would have been in a range of alternative scenarios.

This is a subtle and suggestive idea. I am open to a careful consideration of this sort of strategy. I am however worried that similar individuation worries will afflict Ginet’s suggested approach. That is, will not one have to distinguish *the agent as he actually was* from *the agent as he would have been* (in a range of alternative scenarios), and, thus, will not one have to have some sort of criteria of individuation by reference to which one could say what the relevant “kind” is – “the agent as he actually is”? Will not analogous problems emerge in seeking to specify this kind?³⁰

Ginet additionally finds fault with our way of handling moral responsibility for the consequences of behavior. Ginet focuses on moral responsibility for consequence-universals (or, in an alternative

²⁸ Fischer, *The Metaphysics of Free Will*.

²⁹ Fischer, *The Metaphysics of Free Will*.

³⁰ For the suggestion that the difficulties concerning mechanism-individuation are not insuperable or fatal to the overall Fischer/Ravizza approach, see John Martin Fischer, “Responsibility and Manipulation,” *The Journal of Ethics* 8 (2004), pp. 145–177.

way of putting it, for the obtaining of certain states of affairs). As a card-carrying PAP-ist (believer in the “infallibility” of the PAP!), Ginet disagrees with the intuitive judgment that in such cases as “Joint Assassins” and “Joint Assassins 2” Sam is morally responsible for the Mayor’s dying when he did or shortly thereafter. As Ginet puts it, “[This is] because he could not, by any alternative action or inaction open to him, have avoided its being the case that the Mayor died then.”³¹ He goes on to say:

But, we will hasten to point out, he could have avoided its being the case that the Mayor’s dying occurred *as a result of his action*, and it is his responsibility for *that* state of affairs that vindicates our strong intuition that in these cases Sam is accountable for a nefarious deed. There is a consequence of his action that he could have avoided for which we can hold him responsible, namely, *his* bullet’s striking the mayor in the way that it did. This consequence was, and was intended by Sam to be, causally sufficient for the Mayor’s death. And his being morally responsible for this makes him just as morally reprehensible, and for just the same reasons, as he would have been had there been no overdetermining or preempted cause of the Mayor’s death.³²

In reply, I would first highlight the fact that Ginet has helpfully pointed to the fact that at some level we are in agreement. There are various different aspects of an account of moral responsibility. These include the “content” of moral responsibility – what exactly the agent is deemed morally responsible for. But there are many other facets, including the reasons for holding the agent morally responsible, the degree to which the agent can be praised or blamed, or rewarded or punished, and so forth. Theorists may be in agreement about the latter, even when they disagree about the content of the agent’s moral responsibility. This may seem to diminish the interest of the debates, and in fact I think it does show that there may be no difference *in practice* between apparently competing approaches.

I do however believe that it is important to “get it right” with respect to the content of moral responsibility. Here I prefer what would appear to be a slightly different dialectical strategy. That is, it seems to me that Ginet here may be relying on his allegiance to PAP, and it is at least in part due to his commitment to this principle that he insists that Sam is not (cannot be) morally responsible for the Mayor’s death (in the cases in question). But I should have thought

³¹ Ginet, “Working With Fischer and Ravizza’s Account of Moral Responsibility,” this issue of *The Journal of Ethics*.

³² Ginet, “Working With Fischer and Ravizza’s Account of Moral Responsibility,” this issue of *The Journal of Ethics*.

that it would not be appropriate to rely on a prior inclination to accept PAP in evaluating such cases. I would suggest that we ought to prescind from, or bracket, such theoretical commitments, and seek to evaluate the examples in their own right, as it were. After all, such examples have been invoked to (putatively at least) call into question PAP.³³

I believe that a reasonable and fair-minded person, not already committed to a theoretical principle such as PAP, might well think that Sam is morally responsible for the Mayor's death (in the cases under consideration). I concede that I do not have a knockdown argument that one *must* say this. Rather, I find it plausible, and I do not find any strong reason to reject this natural supposition, apart from any prior commitment to a principle such as PAP. I have thus sought to provide a framework for explaining and justifying this view – a theory according to which Sam exhibits guidance control of the Mayor's death. That is, I have sought to show how such a view is possible, not how it is necessary.

Finally, while I agree that each competing view in this domain has its implausible implications, I wish to point to what I take to be some oddities of Ginet's view here. On Ginet's view, although Sam is morally responsible (in the relevant cases) for the Mayor's dying as a result of Sam's action, he is *not* morally responsible for the Mayor's dying (roughly when he dies). I find it at least a bit odd that, on Ginet's view, Sam is morally responsible for something's happening in a particular way (or as a result of a particular cause), and yet not for that thing's happening.

Of course one cannot be considered morally responsible for everything entailed by what one is morally responsible for, lest one be considered morally responsible for the fact that two plus two equals four or that bachelors are unmarried or that water is H₂O, and so forth. Presumably, it is closer to the truth to say that one can be held morally responsible for any contingent truth (or the obtaining of any contingent state of affairs) non-trivially or "relevantly" entailed by something else for which one is morally responsible. It is hard to specify the pertinent notion of entailment, and I won't be so foolhardy as to try (here). Note, however, that anyone who

³³ I suppose that my co-author and friend, Mark Ravizza, would not now have any problems with the doctrine of PAPal infallibility, insofar as he has been ordained as a Jesuit priest! But I myself resist PAPism of any sort (even though, as a married man, I have taken a vow rather like that of a Jesuit – especially the obedience and poverty part!).

understands that the Mayor has died as a result of Sam's action would thereby understand that the Mayor has died; in this rough sense, "the Mayor's dying" is "contained in" or "entailed (in a special sense) by" "the Mayor's dying as a result of Sam's action. Although of course I have not provided any sort of knockdown argument, I find it implausible that an agent could be morally responsible for its being the case that *P*, but *not* for *Q*, when *Q* is "entailed by" or "included in" *P* (in the indicated sense).

Recall that in "Sharks"

John is on a beach and sees a child struggling in the water. He believes that with very little effort he could save the child from drowning. But, being disinclined to expend any energy to help anyone else, he decides not to save the child and continues his walk along the beach. Unbeknownst to John, a patrol of sharks infests the water between John and the struggling child and would have eaten John, had he jumped in.³⁴

In "Penned-In Sharks," originally suggested to me by David Kaplan, everything is like in "Sharks" except that "the sharks are not swimming freely but are penned in by a man who wants to make sure that the child is not saved; this man would release the sharks if and only if he were to see John jump into the water."³⁵

Ravizza and I argued that in "Penned-In Sharks," but not "Sharks," John is morally responsible for not saving the child. This is because in "Penned-In Sharks," but not "Sharks," John exhibits the requisite sort of control of the child's not being saved. Ginet disagrees with our intuitions about the cases and also finds our attempt at a theoretical explanation of the intuitions unsuccessful. Ginet says:

It is Fischer's and Ravizza's intuition that in "Sharks" John is *not* morally responsible for the fact that the child is not saved by him, because he lacks the required control over that consequence, and I agree. It is, however, also their intuition that in "Penned-In Sharks" John *is* morally responsible, and with this I do not agree. They explain their intuition by claiming that in the alternative scenario where John jumps in the water there occurs a "triggering" event (the bad man's releasing the sharks) that prevents John's action from leading to his saving the child. Now my strong intuition is that John is no more morally responsible for not saving the child in "Penned-In sharks" than he is in "Sharks", that there is no difference between the two cases with respect to his control of that consequence.³⁶

³⁴ Fischer and Ravizza, *Responsibility and Control*, p. 128.

³⁵ Fischer and Ravizza, *Responsibility and Control*, p. 138. Also see Ginet, "Working With Fischer and Ravizza's Account of Moral Responsibility," this issue of *The Journal of Ethics*.

³⁶ Ginet, "Working With Fischer and Ravizza's Account of Moral Responsibility," this issue of *The Journal of Ethics*.

Ginet goes on to argue that our notion of “triggering event” cannot be applied in a way that is consistent with the pattern of intuitive judgments we wish to defend.

Again, this is a very nice challenge. Quite apart from the issue of whether the notion of “triggering event” does the trick (to which I hope to return in future work), here I wish to point out that there are two different senses in which we might be thought to have “explained” our intuition about “Penned-In Sharks.” It is true that we seek to explain the intuition in terms of our theory, which involves the notion of a “triggering event.” But Ginet does not mention another sense in which we sought to “explain” the intuition:

We admit that such cases are puzzling and difficult. But we maintain that this distinction is, upon reflection, *justified*. In ‘Penned-In sharks,’ one holds fixed the actualized conditions, and ‘subtracts’ or disregards the conditions that *would have obtained* in the alternative sequence. And note that this is precisely what one is doing in the Frankfurt-type omissions cases. That is, in the Frankfurt-type “Sloth” case, one is holding fixed the *actual* kind of mechanism, and subtracting off or disregarding the irresistible urges (which occur only in the *alternative scenario*). We agree with such philosophers as Frankfurt, Clarke, and McIntyre and about the Frankfurt-type omissions cases. And if this way of treating such cases is indeed correct, then we submit that our treatment of “Penned-In Sharks” is *also* correct. That is, it is appropriate to treat counterfactual changes in the *second stage* [the path from bodily movement to event in the external world] just like counterfactual changes in the *first stage* [the inner path to the bodily movement].³⁷

Thus, our initial “explanation” of the intuitive difference between “Sharks” and “Penned-In Sharks” was to admit that these are delicate and difficult matters, but to point out that if one says what we believe one ought to say about a whole range of omissions-cases (“Frankfurt-type omissions cases”), then one must also say that John is morally responsible for not saving the child in “Penned-In Sharks.” Our explanation of the intuitive difference between “Penned-In Sharks” and “Sharks” was (in part) in terms of consistency with other cases; we then sought to give a theoretical account, in terms of guidance control, of the configuration of intuitive judgments. Ginet

³⁷ Fischer and Ravizza, *Responsibility and Control*, pp. 138–139. The Frankfurt-type omissions cases are discussed in some detail in Fischer and Ravizza, *Responsibility and Control*, pp. 124–131. In “Frankfurt-type Sloth,” an agent sees a child drowning, believes he could easily save the child, but decides not to bother (for reasons of his own). Unbeknownst to him, some Frankfurt-style counterfactual intervener was present and ensured his slothful choice; that is, even if he had been about to choose to jump into the water, he would have been required to choose the slothful course anyway.

has challenged this latter explanation, but *not* the initial explanation we offered for the asymmetry between “Sharks” and “Penned-In Sharks.”

Now in fairness it should be pointed out that Ginet would presumably disagree with the intuitions we relied upon in analyzing the “Frankfurt-type omissions cases.” As a strict PAPist, Ginet would presumably reject the notion that an agent could be morally responsible for not doing *X*, in a context in which he could not have done *X*. But I wish simply to offer a reminder that we “explained” the alleged intuitive difference between “Sharks” and “Penned-In Sharks” in terms of an allegedly desirable fit with intuitive judgments about other cases. To challenge this claim, one would need to challenge the intuitive judgments about the other cases, or the putative parallel between “Penned-In Sharks” and those cases. I would not be inclined to be sanguine about the success of such a challenge.

6. SOME CONCLUDING REFLECTIONS

Again, I wish to thank Rowe, Mele, Haji, and Ginet for their thoughtful, generous, and challenging papers; I regret that I have been able merely to scratch the surface. I hope to address some of the additional criticisms in the papers published here (and elsewhere) in future work.

Additionally, I hope in future work to help to make some progress toward defending the idea that we can have the sort of control associated with moral responsibility, even in a causally indeterministic world. Although the thrust of my work has focused on causal determinism (as well as God’s omniscience), I also believe that moral responsibility is fully compatible with indeterminism. After all, I believe that our most fundamental attitudes toward ourselves as persons and morally responsible agents should not “hang on a thread”; the discovery that a certain sort of causal indeterminism is true should not shake our confidence in ourselves as persons and morally responsible agents any more than the discovery that causal determinism is true. Our views of ourselves as deeply different from most non-human animals, as setters of ends through deliberation and practical reasoning and as fully and robustly morally responsible, should be resilient to such abstruse cosmological discoveries of the theoretical physicists. And such theories in physics should not be

rejected because they do not fit with our conception of ourselves! It would be bizarre and unacceptable to reject a well-developed and tested theory in physics because it does not comport well with one's views about agency and moral responsibility (or God, for that matter).³⁸

I thus seek to defend a kind of Supercompatibilism. Supercompatibilism is the doctrine that control (of the pertinent kind) and thus moral responsibility are compatible with *both* causal determinism and indeterminism. Note that (as Ginet has essentially pointed out) the fundamental core of the account of guidance control I have offered can be adjusted so as to fit with incompatibilism. I believe that there is nothing in this core that would rule out a non-causal view about the relationship between reasons and actions, or agent-causation. It thus could be considered a template for Supercompatibilism.

If Semicompatibilistic Supercompatibilism is a revolution, it is a gentle revolution. It seeks to show how certain views are possible and defensible, but it does not purport to establish them as beyond reasonable dispute. Mark Twain, with whom I began, once expressed the thought that the notices of his demise were premature. I fervently hope that even those who think Semicompatibilism has no future will have learned some helpful lessons from it. And allow me to point out that when China's great revolutionary, Chairman Mao, was asked what he thought about the French Revolution, he replied, "It is too early to tell."

7. AFTERWORD

About a dozen years ago, I began my monograph, *The Metaphysics of Free Will: An Essay on Control*, with the following quotation from Michael Ross:

Each murder was a fluke – at least that's what I told myself. I knew that I was a 'good' person, that I tried to help people, and certainly I didn't want to hurt anybody. ... Even now, I know that I have done it and know that I could do it again, but I can't imagine myself actually doing it, or even wanting to do it ...

For a long time I looked for excuses. ... But the end result was the same, each murder was a fluke. I made myself believe that there was an excuse and that it

³⁸ For this reason I have always been puzzled by religious persons who resist the attempt to construct a compatibilistic Theodicy; it seems to me that such a person should welcome this sort of reconciliation, lest he be vulnerable to being in an extremely uncomfortable dialectical position in the future.

would never happen again. And the contradiction that it did happen again, and again, was ignored because it didn't fit in with my perception of myself.

I couldn't acknowledge the monster that was inside. ... Sometimes I feel that I am slipping away and I'm afraid of losing control. If you are in control you can handle anything but if you lose control you are nothing.³⁹

Michael Ross was described as a "mild-mannered Cornell graduate" by one author, and as a "scrawny Ivy Leaguer" by some of the residents of Connecticut, into whose houses Ross came as an insurance salesperson. I would hope that, although I could also be truly described in these ways, the similarities stop there; for instance, I am no insurance salesperson. (But there is also perhaps this similarity: Michael Ross spent many years on Connecticut's Death Row, and I was a "junior faculty" person in the philosophy department at Yale!)

Ross struggled to understand his putatively uncontrollable urges, and some of his struggles are chronicled in various essays and also in an online journal. He believed he had some sort of chemical imbalance that resulted in a particular form of psychopathy; this belief was confirmed by certain psychiatrists. Nevertheless, he wondered whether he could have refrained from acting in accordance with his deadly urges. In a striking passage, he writes:

One of my doctors once told me that I am, in a sense, also a victim – a victim of an affliction that no one would want. And sometimes I do feel like a victim, but at the same time I feel guilty and get angry for thinking that way. How dare I consider myself a victim when the real victims are dead? How dare I consider myself a victim when the families of my true victims have to live day by day with the pain of the loss I caused?

So what if it is an affliction? So what if I was really sick? Does that really make any difference? Does that absolve me of my responsibility for the deaths of eight totally innocent women? Does it make the women any less dead? Does it ease the pain of their families? No!⁴⁰

On May 13, 2005, Michael Ross was executed in Connecticut.

On July 22, 2005, Erin Runyon, the mother of the child, Samantha Runyon, murdered viciously by Alejandro Avila, addressed the court

³⁹ Karen Clarke, "Life on Death Row," *Connecticut Magazine* 53 (1990), pp. 51–55; 63–67. For further discussion, see John Martin Fischer and Mark Ravizza, *The Metaphysics of Free Will* pp. 219–220.

⁴⁰ Michael Ross, "It's Time for Me to Die: An Inside Look at Death Row," *The Journal of Psychiatry and Law* 26 (1998), pp. 475–491.

in the sentencing phase of Avila's trial, as reported by *The Orange County Register*:

I have written and re-written what I would say today to you. Part of me doesn't want to speak to you or acknowledge you in any way, but I've decided that I have to address you because I hope to never see you again. I never want to hear your name or see your face. You don't deserve a place in my family's history. And so I want you to live. I want you to disappear into the abyss of a lifetime in prison where no one will remember you, no one will pray for you, no one will care when you die. Since Samantha's death, I have felt more hate and rage than I ever thought possible, but I love that little girl so much that it would be a horrible insult to her to let my hate for you to take more space in my heart and head than my love for her.

I am supposed to speak to the impact of this crime on my life. There is no describing the impact, and I am not sure you're intelligent enough to ever comprehend it anyway. I wrote this statement on the third anniversary of the night you took my baby and hurt her and scared her and crushed her until her heart stopped. She fought. I know she fought. I know she looked at you with those amazing, sparkling brown eyes and you still wanted to kill her. I don't understand it. I never will.

It's like you never learned to think. You have absolutely no concept of how heinous, how egregious your crimes were. I can't help but wonder how it is you survived as long as you did being so stupid.

You killed a child with a loving and passionate heart. Samantha was outrageously bright and funny. She wasn't demanding, she didn't ask for everything under the sun, just to play and have fun as much as humanly possible. Why would you want to take that away? I have researched and really thought about pedophiles and your psychology and blah, blah, blah ... you're a human being, you've known pain and fear. ... Did you pretend that she wasn't real?

I want an apology. Someday I want you to feel the impact of what you did to Samantha. I want you to realize how much you stole. I have to take family photos and my little girl isn't there; she will always be missing. Every happy moment of my life has a moment of gut-wrenching agony because she's not there. And I have to stop and acknowledge how much it hurts to live without her.

Samantha made me feel like I had a purpose on this planet. She was so incredible that I felt sure that if I just did what I could to give her every opportunity to become the best person she could be, and I didn't mess her up in the meantime, she would have done something truly wonderful for this world. She wanted to be a dancer, a teacher and a mother. She was a wonderful storyteller and she wrote all the time. Who knows what she would have become?

But you just don't care. You have no idea of what it is to love someone ... you have no concept of what life is about and yet you were so arrogant as to think you had a right to take it.

For me and my family – our lives were shattered. For the past three years we've been trying to paste it back together, but there's this huge void and the lack of her laughter, of art on the walls of her dancing and singing and running and jumping and swinging and smiling – the lack of Samantha is actually a part of our life now. The pain is impossible to describe, the guilt I feel for bringing that sweet baby into the world only to be tortured and terrified ... I am so sorry I let her down.

And you should be sorry you took her away. You should be so sorry. Not sorry you got caught; not sorry that your wasted life will be taken, as if its worth could ever compare, but sorry that you took a life – the life of a very special little girl.

While everything in me wants to hurt you in every possible way, when I'm very honest with myself what I want more than anything is I want you to feel remorse. Everyone feels alone in our pain and confusion. There is so much misery built into being a human being that I can't fathom what would make you want to add to it.

In choosing to destroy Samantha's life you chose this ... You chose to waste your life to satisfy a selfish and sick desire. You knew it was wrong, but you chose not to think about it. Now you have a lot of time to think about it. Don't waste it. Write it down so that the rest of us might learn how to stop you people. You are a disgrace to the human race.⁴¹

One of the psychologists (Mendel) who testified on behalf of the defense struggled with the notoriously difficult set of questions about genetics, early childhood experience, and control as follows:

Mendel said he did not interview Alejandro Avila, 30, nor was asked to offer a diagnosis of the man convicted of kidnapping, sexually abusing and killing the 5-year-old Stanton girl.

But after reviewing materials provided by the defense, Mendel said he saw "very profound" patterns of alcoholism in the paternal lineage of Avila's family, "quite severe" physical abuse, and sexual abuse of the children that include brothers, a sister and cousins.

Mendel testified that boys suffer more from child sexual abuse. It compromises – with feelings of helplessness – the notion that males are supposed to be able to take care of themselves, and they fear that some feminine aspect inside themselves has brought it on, he said.

"Those are profound, often permanent effects on male victims," Mendel said. "Tragically, that seems to have happened in this case – that sense of doubt, or fear, that one is gay. "Boys tend to feel much more isolated, stigmatized, different and strange," he added.

Asked by Assistant District Attorney David Brent how many boys are sexually molested, Mendel said the estimates are one in six or one in eight.

⁴¹ (http://www.ocregister.com/ocr/2005/07/22/sections/breaking_news/article_607495.php). Accessed on 25 July 2005.

About 30 percent of boys physically abused go on to inflict physical abuse as men, and about half that number of sexually abused boys go on to sexually abuse others, he said.

“(Despite) the fact that terrible things happened to you, you’re not forced to be a molester?” Brent asked rhetorically.

“No, definitely not,” Mendel said.

When further questioned by Assistant Public Defender Denise Gragg, Mendel said how one turns out is the outcome of a “constellation” of factors, including the temperament and resilience of the child victim.

“One of those factors is a person has to choose to be a molester. He’s not forced?” Gragg asked.

“I believe strongly in free will and personal responsibility,” Mendel said. “People have a choice. But I also believe background, genetics and environment put forth enormous pressure.”

In some instances, “it wouldn’t feel like I had a choice. To talk clearly about choice, it’s not as simple as choosing to go to the movies or a ball game,” he said, adding that people “are deeply driven by a lifetime of experiences.”

Asked by Gragg if an attraction to children is not necessarily based on sex, Mendel said it “involves a whole lot of things,” which can include anger, issues of fixation at a certain level of development and feelings by offenders that they are with those “they feel they relate closest to.”

Mendel said many offenders cannot just say, “Ok, it’s bad, I’m not going to do it.”

Also testifying was Avila’s cousin, who gave her name only as Angelica C. She said she was sexually abused by her father, who is Avila’s uncle. She also described the physical abuse that Avila’s father inflicted on the son.

She said her father, Avila’s father and others would get together at family gatherings. She said they “got drunk, knocked each other around, called each other names, (and there) was punching, hitting, yelling, throwing stuff, grabbing chairs. We’d run to the bedroom or to the front of the yard ... we were afraid.” She said Avila’s father “was very hard” on his children, “just mean.” She described physical abuse that included grabbing his sons by the arm, shaking them, hitting them and grabbing a belt and taking them into a room.

Wednesday, Erin Runnion described the pain and loss caused by the “cruel” murder of her child, but also recalled her daughter’s fighting spirit.

Avila was convicted last week of kidnapping the girl on July 15, 2002, from outside her family’s condominium in Stanton, then sexually assaulting and murdering her.

Jurors must now decide whether to recommend that Avila receive the death penalty or life in prison without possibility of parole.⁴²

On my view, moral responsibility is an essentially historical notion – whether one is morally responsible depends crucially on the history of one's choice and behavior. Although I have sought to offer the rudiments of a philosophical framework for understanding the phenomena related to moral responsibility, I will be the first (or, at least among the first!) to admit that the framework is, at best, sketchy and suggestive. It is indisputably vague and programmatic at crucial points. One would like to say something more specific, and more helpful, about cases such as those of Michael Ross and Alejandro Vila and a whole range of other cases where, as it were, the metaphysical rubber meets the moral and judicial road.

I have the (perhaps pathetically deluded) hope that my work can help to structure future research on this difficult but important set of topics. I am convinced that intellectual progress will be made only by being intellectually open-minded. That is, progress will not be made by dismissing traditional philosophical work in these areas; equally, it will not be helpful to dismiss the emerging fields of cognitive science and neuroscience. Indeed, I believe that the path forward should involve an eclectic and broad mix of traditional philosophical analysis and empirical work in neuropsychology, cognitive science, and genetics. Given developments in all of these areas and the potential synergisms, it is an extraordinarily exciting time to be struggling with these great issues!

Department of Philosophy
University of California-Riverside
Riverside, CA
USA
E-mail: john.fischer@ucr.edu

⁴² (http://www.ocregister.com/ocr/2005/05/05/sections/breaking_news/article_508767.php). Accessed on 25 July 2005.