Recent Work on Moral Responsibility*

John Martin Fischer

The topics of moral responsibility and free will have attracted considerable attention recently. In this article I discuss some of this literature. Perhaps inevitably I am selective, and I wish to admit "up front" that the proportionate level of attention I give to various topics unabashedly reflects my interests. This will be then a somewhat opinionated and eclectic (but I hope not entirely eccentric) survey of recent literature on moral responsibility—with an emphasis on the relationship between moral responsibility and free will.

In Section I, I discuss the concept of moral responsibility. In the following sections I turn to the conditions of its application. In Section II, I discuss various approaches to moral responsibility according to which such responsibility requires alternative possibilities. In Section III, I present various strategies for calling into question the traditional association of moral responsibility with alternative possibilities; I focus primarily (although not exclusively) on recent work on the Frankfurt-type examples. In Section IV, I chart out various accounts of moral responsibility which are "actual-sequence" approaches: they reject the requirement of alternative possibilities.

I. THE CONCEPT OF MORAL RESPONSIBILITY

A. *The Strawsonian Account*

One of the most influential accounts of the concept of moral responsibility is the "Strawsonian" account (Strawson 1962). P. F. Strawson argues that we can understand moral responsibility in terms of certain social practices. On this approach, when members of a given society regard someone as a responsible agent, they react to the person (or deem it fitting to react to him) with a characteristic set of feelings and attitudes—for example, gratitude, indignation, resentment, love, respect, and forgiveness. Strawson uses the term "reactive attitudes" to refer to this range of attitudes that "belong to [our] involvement or participa-

tion with others in interpersonal human relationships" (Strawson 1962, p. 194).

Our attitudes toward other persons seem to be importantly different from those we take toward nonhuman animals and inanimate things. That we take the reactive attitudes toward other persons (as opposed to nonpersons) shows that we are engaged with persons in a distinctive way. Nonpersons can be used, exploited, manipulated, or perhaps just enjoyed. But we do not have the reactive attitudes (such as resentment or gratitude) toward them. In contrast to our attitudes toward persons, we view nonpersons from a more detached and "objective" perspective. A broadly Strawsonian approach to moral responsibility analyzes responsibility in terms of the reactive attitudes and certain associated practices, such as punishment and moral reward.[1]

The most sustained, detailed, and nuanced discussion of Strawson's views about the concept of moral responsibility, together with an elaboration of a Strawsonian conception of responsibility, is by R. Jay Wallace (1994). Wallace distinguishes between holding someone morally responsible and that individual's being morally responsible; he then defines being morally responsible (partly) in terms of holding someone morally responsible.

Wallace contends that holding people morally responsible involves "being susceptible to the reactive attitudes" in dealing with them. Whereas P. F. Strawson—and various others in the Strawsonian tradition—think of these attitudes as the wide array of emotions characteristically present in interpersonal relations (as opposed to our relations with nonpersons), Wallace restricts the list to resentment, indignation, and guilt. Wallace contends that his narrower construal of the reactive attitudes is preferable to the wider Strawsonian construal because only on the narrower construal can one understand how these emotions "hang together as a class."[2]

More specifically, Wallace argues that the attitudes of resentment, indignation, and guilt are linked by related propositional objects. Episodes of guilt, resentment, and indignation are all caused by the belief that a moral expectation to which one holds a person has been breached: "The reactive attitudes are explained exclusively by beliefs

1. There is an enlightening discussion of P. F. Strawson's views in the context of a larger development of a Humean account of moral responsibility in Russell (1995). For additional discussion of Strawson's views and "Strawsonian" accounts of moral responsibility, see Bennett (1980); Benson (1990); Watson (1987*b*); McKenna (1998*b*); Galen Strawson (1986, esp. pp. 84–92); Wolf (1981); Fischer (1994, esp. pp. 211–13); and Fischer and Ravizza (1998, esp. pp. 1–8).

2. Wallace (1994), p. 12. Although Wallace makes a case for the narrower construal, I am not entirely convinced that it is necessary to find some additional way in which the attitudes "hang together as a class"—apart from the fact that they characterize what we typically understand to be constitutive of our relations to persons rather than nonpersons.

about the violation of moral obligations (construed as strict prohibitions or requirements), whereas other moral sentiments are explained by beliefs about the various modalities of moral value" (Wallace 1994, p. 38). Not only does the stance of holding someone morally responsible include susceptibility to the reactive emotions, but it also typically involves the application of moral sanctions which serve to express these emotions.

Now Wallace defines "being responsible" in terms of "holding responsible" roughly as follows. An agent is morally responsible insofar as it is fair to hold him morally responsible. This is then a "normative" conception of moral responsibility to the extent that normative issues concerning the fairness of the adoption of the stance of holding someone morally responsible (and thus applying the reactive attitudes and associated sanctions) help to determine whether someone is morally responsible.[3]

B. Oshana's "Accountability" View

The Strawsonian approach stresses what might be called the "social" dimension of moral responsibility—the fact that holding someone (else) morally responsible involves deeming that person a fitting target of certain attitudes and practices. Marina Oshana has offered an alternative "social" conception of moral responsibility (Oshana 1997). Oshana contends that "when we say a person is morally responsible for something, we are essentially saying that the person did or caused some act (or exhibited some trait of character) for which it is fitting that she give an account" (1997, p. 77). On Oshana's approach, being morally responsible for something entails being accountable for it, and "'X is accountable for Y' can be unpacked as 'It is appropriate that X explain her intentions in doing (or being) Y'" (1997, p. 56). According to Oshana, the accountability idea is more basic than the Strawsonian idea of being a suitable target of the reactive attitudes. On her approach, an agent is not morally responsible because she is an apt target for the reactive attitudes; rather, the agent is an apt target for the reactive attitudes because she is morally responsible, that is, it is fitting to demand that she give a certain sort of explanation of her behavior (or her being a particular way).

---

3. Various approaches to free will and moral responsibility have been dubbed "normative" as opposed to "purely metaphysical." Whereas it is useful to have a general rubric for these approaches, they should nevertheless not be assimilated in all respects. Susan Wolf's view is a different sort of normative view. In her stimulating and provocative book, Wolf argues for the "reason view," which is in brief form the claim that "the freedom necessary for responsibility consists in the ability (or freedom) to do the right thing for the right reasons . . . to choose and to act in accordance with the True and the Good" (Wolf 1990, p. 94). For an insightful discussion of the Lockean history of normative approaches and the development of a related view, see Yaffe (1999*a*, in press).

C. Watson's "Two Faces" of Responsibility

In an insightful article, "Two Faces of Responsibility," Gary Watson distinguishes two conceptions—or perhaps aspects—of moral responsibility (Watson 1996). The first notion of responsibility identified by Watson is the "self-disclosure" view. On this view, an agent is morally responsible insofar as he has the capacity to choose ends freely and act in accordance with such choices. An agent with such a capacity can be considered morally and not merely causally or "superficially" responsible for his behavior, according to Watson. Watson says, "moral accountability is only part, and not necessarily the most important part, of our idea of responsibility. The self-disclosure view describes a core notion of responsibility that is central to ethical life and ethical appraisal. In virtue of the capacities identified by the self-disclosure view, conduct can be attributable or imputable to an individual as its agent and is open to appraisal that is therefore appraisal of the individual as an adopter of ends. Attributability in this sense *is* a kind of responsibility" (1996, p. 229). Watson elaborates, "attributability has an importance to ethical life that is distinct from concerns about accountability. Responsibility is important to issues about what it is to lead a life, indeed about what it is to have a life in the biographical sense, and about the quality and character of that life. These issues reflect one face of responsibility (what I call its *aretaic* face). Concerns about accountability reflect another" (1996, p. 229).

Watson's self-disclosure conception of moral responsibility corresponds to the idea of "deep attributability" or (say) an action's being "really an agent's" or "the agent's own" in some appropriate sense. Various philosophers think of moral responsibility in this way. For example, Derk Pereboom says, "for an agent to be morally responsible for an action is just for the action *really to belong to the agent.* Equivalently, but in classical phrasing, for an agent to be morally responsible for an action is for it to be *imputable* to the agent."[4]

Watson's self-disclosure or real imputability notion of moral responsibility is different in an interesting way from both "social" conceptions of responsibility discussed above: the Strawsonian approach and Oshana's accountability approach. Watson says:

> If someone betrays her ideals by choosing a dull but secure occupation in favor of a riskier but potentially more enriching one, or endangers something of deep importance to her life for trivial ends (by sleeping too little and drinking too much before important per-

---

4. Pereboom (1999, p. 10). Interestingly, Pereboom brings Watson's two conceptions of responsibility together in the following alternative formulation: "For an agent to be morally responsible for an action is for this action to belong to the agent in such a way that she would deserve blame if the action were morally wrong, and she would deserve credit or perhaps praise if it were morally exemplary" (ibid., p. 10).

formances, for example), then she has acted badly—cowardly, self-indulgently, at least unwisely. But by these assessments we are not thereby *holding* her responsible, as distinct from holding her to be responsible.[5] To do that, we would have to think that she is accountable to us or to others, whereas in many cases we suppose that such behavior is "nobody's business." Unless we think she is responsible to us or to others to live the best life she can—and that is a moral question—we do not think she is accountable here. If her timid or foolish behavior also harms others, and thereby violates requirements of interpersonal relations, that is a different matter. (1996, p. 231)

Watson employs the term 'accountability' to refer to a more "social" notion of moral responsibility which includes the reactive attitudes and associated practices, that is, roughly a Strawsonian notion of moral responsibility. (Note then that Watson's accountability conception of responsibility is different from Oshana's.) He argues that the distinction between the two kinds of responsibility helps to explain our ambivalence (when we are indeed ambivalent) toward the vicious criminal who is himself a victim of an abusive childhood.[6] On the one hand, we tend to hold such an individual responsible in the aretaic sense because "his conduct is attributable to him as an exercise of his 'moral capacities'. It expresses and constitutes his practical identity, what he stands for, what he has made of his life as he found it" (Watson 1996, p. 240). On the other hand, there is an inclination to doubt that such a person can legitimately be held morally responsible. As Watson puts it, "This ambivalence mirrors the two faces of responsibility. . . . Facts about his formative years give rise to the thought that the individual has already suffered too much and that we too would probably have been morally ruined by such a childhood. What is inhibited by these concerns is accountability blame" (Watson 1996, p. 240).

Thus Watson explains our natural ambivalence about criminals such as Robert Alton Harris in terms of the conflict between the two faces of responsibility.[7] He concludes, "underlying these distinguishable perspectives are two sets of overlapping interests, both central to the ethical life. One set of interests hinges on our concern with living a good human life, with models and ideals of human possibility. The second set of interests pertains to social regulation and (more obscurely) to retributive and compensatory justice. In the end, I doubt that these interests can be fully held apart. But it is important to see, as I have tried to show, that they have distinct sources" (1996, p. 243).

5. Here Watson is obviously not in agreement with Wallace's project of defining being morally responsible in terms of holding one morally responsible.

6. For discussion of such cases, focusing on the infamous Robert Alton Harris, see Watson (1987*b*) and McKenna (1998*b*).

7. Watson discusses Harris at length in Watson (1987*b*).

D. Responsibility and Autonomy

For many years I have been struck by the fact that there are "parallel literatures" which discuss "moral responsibility" and "autonomy." In many ways the literatures are isomorphic. For example, "hierarchical" approaches (involving the apparatus of higher-order mental states— states directed at [say] "first-order" mental states) are employed in the literature on moral responsibility as well as on autonomy.[8] Also, similar debates crop up in both literatures; in both literatures there are debates about whether the relevant notion—responsibility or autonomy—is essentially "historical."[9]

What exactly is the relationship between moral responsibility and autonomy? Are the terms simply two names for the same concept? I believe that autonomy entails moral responsibility, but it is not the case that moral responsibility entails autonomy. It seems to me that one can be morally responsible without meeting the additional requirements of autonomy. For example, a subservient wife may be morally responsible, and yet entirely submissive, taking her cues from her husband in all important respects. She may well meet the more minimal conditions for moral responsibility without having the more robust self-governance required for autonomy.[10] Moral responsibility is then a more abstract, minimal notion; it is the "gateway" to the more substantive notion of autonomy, but it is not in itself sufficient for autonomy.[11]

II. THE ALTERNATIVE-POSSIBILITIES REQUIREMENT FOR RESPONSIBILITY

A. Alternative-Possibilities Control

The broadly-speaking "Aristotelian" conditions on moral responsibility require that an agent meet certain "epistemic" and "freedom-relevant" conditions. Put "negatively," the agent must not be ignorant of certain

---

8. Here Harry Frankfurt's work has been particularly influential. See, e.g., Frankfurt (1971). For critical discussion, see, e.g., Thalberg (1978); Shatz (1985); and Mele (1995, esp. pp. 65–80). On autonomy, see, e.g., Dworkin (1970, 1988); and Christman (1989).

9. In the literature on moral responsibility, see Frankfurt (1975, esp. pp. 121–22); and Fischer and Ravizza (1998, pp. 170–206). In the literature on autonomy, see Christman (1991); and Mele (1995, esp. pp. 156–62).

10. Some helpful discussion of similar examples can be found in Wolf (1987); Friedman (1986); and Oshana (1998).

11. There are helpful and comprehensive recent treatments of autonomy in Mele (1995); and Berofsky (1995). Randolph Clarke in personal correspondence has suggested that it might be possible that one be autonomous without being morally responsible. To be morally responsible, one presumably has to be capable of understanding moral reasons and responding to them suitably. It might be the case, Clarke suggests, that one satisfies the requirements for self-governance, being sensitive to reasons of some kinds, but that one is oblivious to moral reasons.

crucial features or consequences of his behavior, and he must not be "forced" to behave as he does. Both the epistemic and the freedom-relevant conditions are important, but I focus primarily on the freedom-relevant condition.[12]

Traditionally the most influential view about the sort of freedom necessary and sufficient for moral responsibility posits that this sort of freedom involves the availability of genuinely open alternative possibilities at certain key points in one's life. Without this sort of alternative possibility, it is alleged, one is compelled to do as one actually does. Some philosophers talk in terms of freedom; others employ the term, 'control'. I prefer 'control' because it highlights the fact that mere chance occurrences do not secure the satisfaction of the relevant requirement. The traditional view then is that moral responsibility for behavior requires the sort of control that involves genuinely available alternative possibilities at some point suitably related to the time of the behavior in question. We might call this sort of control "alternative-possibilities" control.[13]

The intuitive picture behind the alternative-possibilities control requirement is that moral responsibility requires that the agent select one from among various genuinely open paths the world might take. There are two important ideas here. One is that there must be various paths genuinely available to the agent (at least at some times suitably related to the time of the behavior under consideration). The second idea is that the agent (and not some outside force or mere chance) selects which path will be the path into the future. It seems to me that both ideas are important components of the traditional conception of the sort of control associated with moral responsibility—alternative-possibilities control.

B. The Consequence Argument

Many proponents of the alternative-possibilities control requirement for moral responsibility are incompatibilists about moral responsibility and causal determinism. This is because they are incompatibilists about causal determinism and the existence of the relevant sorts of alternative possibilities. One of the most influential arguments for the incompatibility of causal determinism and alternative-possibilities control is what Peter Van Inwagen has dubbed "the consequence argument."[14]

To develop the consequence argument, I begin with a rough and simplistic account of causal determinism. Causal determinism, for the purposes of the argument, can be defined as the claim that a complete

12. For useful discussions of the epistemic conditions, see Feinberg (1986), esp. pp. 269–315; and Haji (1998), esp. pp. 172–74.

13. Elsewhere I have called it "regulative" control: Fischer (1994); and Fischer and Ravizza (1998).

14. For some formulations and discussions of this argument, see Van Inwagen (1983), esp. pp. 55–105; Ginet (1990), esp. pp. 90–123; and Fischer (1994), esp. pp. 1–110.

statement of the laws of nature and a complete description of the (temporally nonrelational or "genuine") facts about the world at some time T entail every truth about the world after T. Now the argument can be stated in simple form as follows. Consider some choice of mine C which I actually make at time T_2. If causal determinism is true, then the total state of the world at T_1 together with the laws of nature entail that I make C at T_2. Thus, in order for me to have made a different choice at T_2, it must have been the case that the state of the world at T_1 was different from what it actually was, or some law of nature which actually obtained did not obtain. But I cannot at any time so act (or choose) that the past would have been different from the way it actually was. And I cannot at any time so act (or choose) that some actual law of nature would not have obtained. That is, the past and the natural laws are "fixed" in the relevant sense. Thus, it seems that I could not have chosen otherwise at T_2. And the argument can obviously be generalized to show that if causal determinism is true, then I never can choose or behave differently from the way in which I actually choose (or behave). That is, if my choices and actions are consequences of the past and laws of nature, then I do not have genuinely open alternative possibilities (of the sort required for moral responsibility).

Not everyone has accepted the soundness of this argument, and various philosophers would maintain that we can have "genuinely available" (in some relevant sense) alternative possibilities, even in a causally deterministic world.[15] But many are inclined to accept that the argument (in some version or other) establishes that if causal determinism were true, then we would not have the relevant sort of alternative possibilities (of choice or action). Thus, various philosophers who firmly believe that we are morally responsible—and accept the alternative-possibilities control requirement on moral responsibility—have concluded that causal determinism is false. They have then adopted indeterministic models of moral responsibility.

C. The "Rollback" Argument

Return to the intuitive picture behind the alternative-possibilities control requirement on moral responsibility: an agent is envisaged as selecting from among various genuinely available paths into the future. I suggested above that there are two elements to this picture: there must be alternative possibilities, and the path taken must be selected by the agent. It is important to see that there is a tension between these two elements. Embracing indeterminism is a good way to allow for alternative possibilities; it certainly blocks what many philosophers take to be the most potent argument for the lack of alternative possibilities (the

15. Lehrer (1976, 1980, 1997*a*, 1997*b*); Horgan (1979, 1985); Lewis (1981); Kapitan (1996); Vihvelin (1998); Yaffe (in press), pp. 153–70; Bok (1998).

consequence argument). But precisely by allowing for alternative possibilities in this way, one attenuates the claim that it is the agent who selects the path into the future.

When an agent has control over his behavior in the relevant way, the behavior must flow from him in an appropriate manner. Following Timothy O'Connor, I call this sort of control—that involves the behavior's being an "outflowing of the agent" in the right way—"agent control" (O'Connor 1993*b,* esp. p. 500). The traditional conception of the sort of control required for moral responsibility involves both the possession of alternative possibilities and agent control. These two elements can be seen as aspects of the idea that, when an agent acts with free will, he makes a difference to the world through his behavior. When an agent lacks alternative possibilities, he does not appear to make a *difference* to the world. And when he lacks agent control, *he* does not appear to make a difference to the world. Whereas causal determinism is a threat to the first element, the lack of causal determinism is a threat to the second element (agent control).

Imagine that causal determinism is false in such a way as to make our choices and behavior genuinely random. This sort of indeterminism would imply that an individual's behavior is relevantly similar to the output of a truly random device, such as a random number generator. But if so, surely the individual does not control his behavior. He does not select the path the world will take; he does not make a difference to the world through his behavior. It is well known that the introduction of genuine randomness (as an alternative to causal determination) does not help to secure the relevant kind of control.

Now consider another sort of indeterminism—which actually has a good chance of obtaining, according to contemporary physics. This sort of indeterminism is not randomness. Rather, it simply posits that there are residual extremely small but nevertheless nonnegligible possibilities that macroscopic objects will not act as they typically do (that is, as they are supposed to, according to the laws of nature [which are, on this view, probabilistic generalizations that fall slightly short of universal generalizations]). This residual possibility of "weird" behavior of macroscopic objects is putatively supervenient on quantum indeterminacy at the microlevel. Although for practical purposes the residual possibilities can be ignored, their presence issues in the falsity of causal determinism.

Let us suppose that the world is as just described; that is, let us suppose that a certain sort of indeterminism, "almost determinism," obtains. There appears to be a powerful argument from almost determinism to the lack of agent control—the "rollback argument."[16] Imagine, as above, that I make a certain choice C at time T_2, and suppose that the

16. For discussions, see Van Inwagen (1993), esp. pp. 126–50; and Mele (1995), pp. 195–209.

doctrine of almost determinism obtains. Now imagine that we "roll back" the world to its state at T_1, holding everything about the world—including all my values, beliefs, and motivational states—fixed (i.e., making sure that they are just as they actually are at T_1), and we let the world "unfold" into the future an indefinitely large number of times. If we can indeed imagine this, it will turn out that sometimes I will make the choice C (the one I actually made) at T_2 and sometimes I will not make this choice. That sometimes I make choice C and sometimes I do not is disturbing: after all, nothing about me and my motivational states differs between the scenarios in which I make choice C at T_2 and those in which I do not. Given this, it can seem that I do not actually control my behavior in the sense relevant to ascriptions of moral responsibility; that is, it can seem that I lack agent control. (A difference is made to the world, but it is not I who makes it.)

Perhaps the same point can be seen by noting that in a context of indeterminism, we do not have an explanation of why the agent made the particular choice he actually made rather than another. The rollback argument points to the fact that the antecedent conditions—say my standing desires, values, intentions, and plans—cannot in themselves explain why I actually choose C rather than something else (given indeterminism). The sort of explanation which is lacking here is called a "contrastive explanation."[17]

The rollback argument is a major challenge to libertarians (indeterminists who believe we have the sort of control associated with moral responsibility). But the challenge has certainly not gone unanswered in recent years. One of the most persuasive and appealing ways of addressing this challenge is set out by Robert Kane in his comprehensive and forceful defense of libertarianism (1996, esp. pp. 124–90). Put briefly, Kane wishes to distinguish between what might be called "antecedent-ensuring control" and "simultaneous control." Antecedent-ensuring control implies that, given all the relevant antecedent conditions, an agent can ensure that a certain result (a choice, action, omission, or consequence) will ensue. Antecedent-ensuring control is shown by the rollback argument to be inconsistent with indeterminism (and thus almost determinism). But Kane argues that there is a different sort of control that an agent may display when he voluntarily chooses (or behaves) as he does, simultaneous control, and this sort of control is not necessarily ruled out by causal indeterminism. Additionally, an agent in an indeterministic world may have the dual power to exhibit simultaneous control in the actual sequence and also in relevant alternative sequences. Further, according to Kane, moral responsibility does not re-

17. There are discussions in Nagel (1995); Double (1988); Clarke (1996*b*); and Kane (1996), pp. 174–79.

quire antecedent ensuring control, but it does require (dual) simultaneous control.[18]

Similarly, Randolph Clarke contends that antecedent-ensuring control is not necessary for the kind of free will (and control) necessarily linked with moral responsibility (Clarke 1995). Further, Clarke insightfully points out that causal indeterminism does not imply that when an agent chooses (or acts), he lacks control in so choosing (or acting). Rather, it entails that in certain sorts of cases the agent runs the risk of choosing (or acting) without control. (Clearly, the relevant sort of control here is some sort of "simultaneous" control.) Indeterminism (and thus almost determinism) does rule out antecedent-ensuring control, but it does not necessarily undermine simultaneous control, although it does establish an antecedent risk that there will be no simultaneous control. (In the case of almost determinism, this risk is extremely low.)

D. Libertarian Accounts

Thus, although the two elements (the existence of alternative possibilities and the presence of agent control) in the picture behind the traditional conception of the sort of control associated with moral responsibility are in tension, there are various plausible ways of seeking to render the elements harmonious. Indeed, there has been much excellent recent work not only addressing the relationship between indeterminism and agent control but also developing positive accounts of what this sort of control consists in (on a libertarian model).

1. *Ginet's simple indeterminism.*— One might distinguish different strategies for giving a "positive" account of libertarian control.[19] The three most salient such strategies might be dubbed "simple indeterminism," "probabilistic" approaches, and "agent causality."[20] Carl Ginet has argued for a noncausally deterministic relationship between antecedent reasons and the subsequent choices and actions of the agent; following O'Connor, I shall call Ginet's approach, "simple indeterminism."[21] On

18. I believe that Kane would deny that moral responsibility requires contrastive explanation. He would, however, point out that there is another sort of rational explanation associated with moral responsibility, according to which one shows that the agents acted for reasons they wanted to act on when they acted, rather than doing something "by mistake, or accidentally, inadvertently, or against their wills" (Kane 1996, pp. 178–79). Randolph Clarke (ibid.) argues that indeterminism is perfectly compatible with the possibility of providing adequate rational contrastive explanation (and thus causal contrastive explanation).

19. I employ the term 'positive' to distinguish this part of libertarianism from the "negative" component, which is the argument against the compatibility of causal determinism and the relevant sort of control (freedom) and moral responsibility (the "consequence argument").

20. For a useful overview, see the introductory essay in O'Connor (1995).

21. Ginet (1989, 1990).

Ginet's view, there is a simple (causally unstructured) mental action at the core of every causally complex action.[22] In some contexts these simple mental actions are not part of a more complex action; this would be the sort of case in which someone simply mentally says a word or voluntarily forms an image.

But most cases involve the simple mental action as a component of a larger action. On Ginet's view, the simple mental action constitutes the core element of one's voluntarily exerting the body, causally producing that exertion. Ginet requires that in order for the bodily actions to be free, the simple mental actions must be "undetermined," by which he means not nomically necessitated and not deterministically caused. On Ginet's view, another sort of relationship—a noncausal relationship—is sufficient to make it the case that an action is explained by antecedent motivational states of the agent; this relationship is cashed out in terms of the intentional contents of the antecedent motivational states ("reasons"), the intentional content of the action, and the "directly referential" relation between the action and its concurrent intention. Clearly, the main challenge for such a view is to defend the claim that the sort of control associated with moral responsibility is compatible with the lack of causation.[23]

2. *The simple probabilistic model.*—There are various "probabilistic" models of libertarian freedom (and thus moral responsibility). On one approach, which I shall dub the "simple probabilistic model," antecedent motivational states such as desires and beliefs do indeed cause subsequent choices but not via a causally deterministic path. Thus, this approach is committed to the coherence of nondeterministic event causation (i.e., causation in which the relata of the causal relation are both events).

Peter Van Inwagen is a theorist who at least believes that such a model is coherent, although he concedes that he cannot give a full defense of it (Van Inwagen 1983, esp. pp. 146–50). Van Inwagen is quite convinced by the negative component of libertarianism—the consequence argument. He is also very confident that we do indeed have free will and moral responsibility. He thus concludes that some positive model of libertarian (indeterministic) freedom and moral responsibility must be coherent. He does not contend that the probabilistic model is the only coherent libertarian picture, but he does conclude that it is at least a coherent picture. In a colorful passage, Van Inwagen says:

> Now I wish I knew *how it could be* that, for example, our thief had a choice about whether to repent, given that his repenting was caused, but not determined, by his prior inner states, and given

22. For helpful exegesis and discussion, see O'Connor (1996), esp. pp. 143–51.
23. For a critical discussion of Ginet's view, see ibid.; Ginet discusses agent-causation views and defends his approach in Ginet (1997).

that no other prior state "had anything to do with"—save negatively: in virtue of its non-interference with—his act. I have no theory of free action or choice that would explain how this could be. But then neither have I, and neither has the compatibilist, any theory of free action or choice that would explain how any of the propositions in the above list [which appear to entail the incompatibility of causal determinism and alternative possibilities] could be false. Moreover, it is certainly not unheard of in philosophy for an incontrovertible argument to force upon one a puzzling conclusion that one has no theoretical account of. . . . I must choose between the puzzling and the inconceivable. I choose the puzzling. (Van Inwagen 1983, pp. 149–50)

The inconceivable, for Van Inwagen, is that causal determinism should be compatible with the sort of control that involves alternative possibilities, which Van Inwagen deems necessary for moral responsibility. He thus opts for the merely puzzling, which is that the probabilistic model of control is coherent (despite the power of the rollback argument). Van Inwagen has focused his energies mainly on the negative component of libertarianism; in contrast, Robert Kane has attended more to the positive component. He offers a more nuanced probabilistic model, which I shall call the "complex probabilistic model" of libertarian free will and moral responsibility.

3. *Kane's probabilistic account.*—Kane's most fully articulated account of libertarian freedom is presented in *The Significance of Free Will.* On his approach, freedom of the will—or the control relevant to moral responsibility—is directly manifested only in contexts of conscious deliberation.[24] Whereas there are various such contexts, two of the most salient—which are treated essentially alike by Kane—are moral and prudential choices. Kane holds that there is a basic structure in contexts of moral and prudential deliberations.

First, the agent experiences a motivational conflict, feeling inclined to pursue each of two or more incompatible courses of action. The different motivations for each such action are deemed by Kane to be incommensurate. Second, and crucially, such conflicts result in an "effort of will." In a case of moral decision making, it is a struggle to act in conformity with perceived obligation rather than contrary motives. The outcome is uncertain in the agent's mind prior to the moment of choice. Kane claims that an effort of will is an indeterminate process, analogous to the indeterminacy of the position and momentum of an individual particle at the quantum level prior to a measurement. Kane's suggestion is that our brains may (via "chaotic processes"—which are extremely sensitive to small changes in initial conditions) amplify indeterminate events at the microlevel. Third, resolution of the indeterminacy

24. I have been aided in this summary by O'Connor (1996), esp. pp. 152–53.

is brought about through an agent's choice. Because of the indeterminacy of the effort of will which precedes it, the choice is a causally undetermined event.

On Kane's view, the agent's prior character and motives—his standing desires, beliefs, values, plans, and so forth—explain why the agent goes through the process of trying to sort out his ends and purposes (the effort of will). But these prior states do not thereby explain why the actual choice is made (rather than certain others). As I noted above, on Kane's view the agent lacks antecedent ensuring control of his choice, but according to Kane, this is consistent with its being the case both that when he chooses he does so voluntarily (and thus displays simultaneous control), and, had he chosen differently, he would have done so voluntarily (and thus displayed a kind of simultaneous control). We cannot give a contrastive explanation in terms of his antecedent motivational states of his actual choice, but we can say that when he chooses as he actually does, he chooses for a reason.[25]

A third approach to giving a positive libertarian account of the sort of freedom involved in moral responsibility is the "agent-causal" approach. One of the most influential historical developments of the agent-causal picture is in Thomas Reid.[26] In contemporary philosophy this view is frequently associated with Roderick Chisholm and Richard Taylor.[27] Agent causation has received much attention in recent work.[28] Two philosophers have recently developed sophisticated articulations of agent causation, together with defenses against criticisms of the view: Timothy O'Connor (1993*b*, 1995, 1996, in press) and Randolph Clarke (1993, 1996*a*).

4. *O'Connor's agent-causal theory.*—Various agent-causal theorists posit different events as brought about by agent causes: bodily movements, brain events, the agent's endeavoring or undertaking something, an intention, a choice, or a volition. Additionally, as Clarke points out, "sometimes the event that is directly caused by the agent is regarded as the action, or an event that begins the action, while in other cases the agent's causing this event is considered the action. One version holds that the event that the agent directly causes is identical with the agent's causing that very event" (Clarke 1996*a*, p. 20). On O'Connor's approach, the directly agent-caused events are "determinate (immediately

25. For a critical discussion, see ibid., esp. pp. 153–56.
26. Reid (1852). An excellent recent development and discussion of Reid's theory of action is in Rowe (1991).
27. See Chisholm (1966, 1971, 1976*a*, 1976*b*); Taylor (1966, 1992). As Robert Kane has reminded me, it is interesting that both Chisholm and Taylor have subsequently rejected and abandoned the agent-causal approach. See Chisholm (1982*a*, 1982*b*, 1995); and Taylor (1982).
28. See, e.g., Thorp (1980); and Zimmerman (1984).

executive) intentions to act in various ways" (1996, p. 145). O'Connor gives a brief account of the agent-causal relation as follows:

> Wherever the agent-causal relation obtains, the agent bears a *property* or set of properties that is "choice-enabling" (i.e., in virtue of such properties, the agent has a type of causal power which, following Reid, we may term "active power"). But this "active power"— the causal power in virtue of which one has freedom of will—is not characterized by any function from circumstances to effects (as is the case with event-causal powers). For the properties that confer such a capacity do not themselves (in the appropriate circumstances) necessitate or make probable a certain effect. Rather, they (in conjunction with appropriate circumstances) *make possible* the direct, purposive bringing about of an effect *by the agent* who bears them. (1996, p. 145)

On O'Connor's approach, then, agent causation is a species of causation, "production," or "bringing about," but it is a different sort of causation from event causation, "though the very same relation of causation is involved, these properties [the properties that confer the agent-causal capacity] give rise to a fundamentally different type of causal power— one that in suitable circumstances is exercised at will by the agent, rather than of necessity, as with objects that are not self-determining agents" (1996, p. 145).

O'Connor's view is then a kind of "traditional" account of agent causation. Randolph Clarke has pointed out that the traditional account of agent causation has two salient features. First, agent causation is a genuine species of causation, but fundamentally different in nature from event causation. Second, the agent-caused event is not thought also to be subject to event causation.

5. *Clarke's "nontraditional" agent-causal account.*—In contrast to O'Connor's view, Clarke has presented a "nontraditional" account of agent causation, which he calls the "causal agent-causal view" (Clarke 1993, 1996*a*). On this view, agent causation involves exactly the same sort of causation as event causation. The only difference lies in the relata of the causal relation; in the case of event causation, the relata are both events, whereas in the case of agent causation, the first relatum is an agent and the second an event.

A second feature of Clarke's view is quite original and suggestive. He contends that the event that is agent-caused—the agent's action—can also be event-caused. Such an action would be free (in the sense relevant to ascriptions of moral responsibility) only if it is caused indeterministically. On Clarke's approach, an agent acts with free will insofar as she agent-causes her action and her action is indeterministically caused (in an appropriate way) by her reasons. As Clarke puts it, "the agent's acting with free will consists (crucially but not wholly) in her action's being caused, in this way, by her and by her reasons" (Clarke 1996*a*, p. 26).

One of the interesting implications of adopting the nontraditional agent-causal view is that it can reply in a natural way to a vexing puzzle for the traditional approach. It seems that the traditional agent-causal view cannot explain why an agent acts at the time at which she acts. If I raise my hand at a certain time T, it seems that the fact that I act precisely when I do—at T—cannot be explained simply by pointing to a causal relation between an agent (me) and an event (act of raising my hand). But if my act of raising my hand is also caused (indeterministically) by my reasons, then we can in principle explain why I act at the time at which I do act (given that, as Clarke contends, indeterministic reasons explanations can be adequate).

III. MORAL RESPONSIBILITY AND ALTERNATIVE POSSIBILITIES

A. *Wallace's Theory*

I pointed out in the previous section that many philosophers contend that alternative-possibilities control is a necessary condition for moral responsibility. This is the powerful and influential idea that in order legitimately to be held accountable for my actions, I must have had freedom to choose and behave differently at some pertinent point along the path to the actions in question. Various philosophers who embrace this requirement believe that one must deny that causal determinism obtains in order to preserve our view that we are (at least sometimes) genuinely and "deeply" morally responsible agents.

There are, however, potent challenges to the alternative-possibilities requirement. The first such strategy is developed by R. Jay Wallace (1994). As I said in the first section of this article, Wallace presents a sophisticated Strawsonian account of the concept of moral responsibility. In order to elucidate the conditions under which we apply this concept, Wallace considers our practices of excusing agents and also exempting them from moral responsibility. Excusing an agent entails deeming him a morally responsible agent (and thus in principle accountable for his behavior) but not appropriately held responsible in a particular context. Exempting an agent from moral responsibility is, in contrast, a more global judgment: this entails that the agent does not have the general characteristics in virtue of which he can be held accountable for any of his behavior. Wallace's strategy is carefully to analyze our practices of excuse and exemption and to argue that they do not imply that moral responsibility requires the genuine availability of alternative possibilities.

It is tempting indeed to suppose that the various sorts of situations in which we would be inclined to excuse a person can be systematized by the principle that an agent cannot be held morally responsible unless he has alternative possibilities available to him. But Wallace argues that a different explanation of the excuses is available—one that more ade-

quately explains the force of excuses over a broad range of cases. According to this approach, excuses serve to show that an agent has not really done anything morally wrong; that is, they show that an agent has not intentionally violated some moral requirement that we accept (Wallace 1994, p. 127). Surely it is unfair to blame someone if he has not done anything wrong in the first place (Wallace 1994, p. 135).

Wallace also contends that we exempt agents from responsibility when they do not possess certain general capacities: the capacities to grasp moral reasons and control behavior in light of them. These capacities—the powers of "reflective self-control"—give the agent a certain kind of "normative competence."[29] Wallace argues that our inclinations to exempt agents from moral responsibility are best explained by reference to the lack of the powers of reflective self-control.

Thus, Wallace argues that our practices of excuse and exemption do not imply that moral responsibility requires alternative possibilities. Similarly, Wallace argues that these practices do not in any way presuppose the falsity of causal determinism. Wallace's route to these conclusions is via general considerations about the practices that constitute moral responsibility in our culture. Another route to the same conclusion employs a certain thought experiment developed by Harry Frankfurt (1969). This sort of thought experiment was developed by Frankfurt precisely to call into question the Principle of Alternative Possibilities, according to which moral responsibility (for actions) requires alternative possibilities.

B. Frankfurt-Type Examples

Here is a particular version of a "Frankfurt-type case."[30] In this sort of case, a crucial role is played by some kind of involuntary sign or indication of the agent's future choices and behavior.[31] Suppose Jones is in a voting booth deliberating about whether to vote for Gore or Bush. (He has left this decision until the end, much as some restaurant patrons wait until the waiter asks before making a final decision about their meal.) After serious reflection, he chooses to vote for Gore and does vote for Gore by marking his ballot in the normal way. Unbeknownst to him, Black, a liberal neurosurgeon working with the Democratic Party, has implanted a device in Jones's brain which monitors Jones's brain activities.[32] If he is about to choose to vote Democratic, the device simply continues monitoring and does not intervene in the process in any way. If,

29. Wallace (1994), p. 162. For a development of an alternative notion of normative competence, see Wolf (1990).

30. Frankfurt (1969). I present the cases similarly in various places, including Fischer (in press), on which I rely considerably in this article.

31. For this kind of Frankfurt-type case, see Blumenfeld (1971).

32. Of course, this sort of example is a highly implausible science-fiction scenario, since most neurosurgeons are certainly not liberal!

however, Jones is about to choose to vote (say) Republican, the device triggers an intervention which involves electronic stimulation of the brain sufficient to produce a choice to vote for the Democrat (and a subsequent Democratic vote).

How can the device tell whether Jones is about to choose to vote Republican or Democratic? This is where the "prior sign" comes in. If Jones is about to choose at T_2 to vote for Gore at T_3, he shows some involuntary sign—say a neurological pattern in his brain—at T_1. Detecting this, Black's device does not intervene. But if Jones is about to choose at T_2 to vote for Bush at T_3, he shows an involuntary sign—a different neurological pattern in his brain—at T_1. This brain pattern would trigger Black's device to intervene and cause Jones to choose at T_2 to vote for Gore and to vote for Gore at T_3.

Given that the device plays no role in Jones's deliberations and act of voting, it seems to me that Jones acts freely and is morally responsible for voting for Gore. And given the presence of Black's device, it is plausible to think that Jones does not have alternative possibilities with regard to his choice and action.

At this point it may be objected that, despite the initial appearance, Jones does have at least some alternative possibility. Although Jones cannot choose or vote differently, his brain can still exhibit a different neurological pattern $N*$ (from the one he actually exhibits, N). I have called such an alternative possibility a "flicker of freedom." The flicker theorist contends that our moral responsibility always can be traced back to some suitably placed flicker of freedom; our responsibility is grounded in and derives from such alternative possibilities.[33]

It seems that one can always find a flicker of freedom in the Frankfurt-type cases insofar as they are developed as prior-sign cases. That is, the agent will always at least have the power to exhibit an alternative sign. But I contend that the mere involuntary display of some sign—such as a neurological pattern in the brain, a blush, or a furrowed brow—is too thin a reed on which to rest moral responsibility. The power involuntarily to exhibit a different sign seems to me to be insufficiently robust to ground our attributions of moral responsibility.

I have argued for this contention at some length elsewhere (Fischer 1994, pp. 131–59). The debate here is subtle and complex; there are different versions of the flicker strategy and various different responses. But for my purposes in this article perhaps it will be enough to reiterate one line of argument I have developed against the flicker approach. Note that in the alternative sequence (in which Jones shows neurological pattern $N*$, which is indicative of an impending decision to vote for Bush), the sign is entirely involuntary, and the subsequent decision and

---

33. Important early developments of the flicker-of-freedom strategy can be found in Van Inwagen (1978, 1983); and Naylor (1984).

vote are produced electronically. Thus, in the alternative sequence Jones cannot be said to be choosing and acting freely and, similarly, cannot be thought to be morally responsible for his choice and action.

Imagine, just for a moment, that there are absolutely no alternative possibilities, even the flimsy and exiguous flickers of freedom we have recently been entertaining. An alternative-possibilities control theorist would say that under such circumstances the relevant agent cannot be morally responsible for his choice and action. Now add the flickers of freedom we have been considering—the power to exhibit a different neurological pattern, N^*. I find it very hard to see how adding this power can transform a situation in which there is no moral responsibility into one in which there is moral responsibility. How can adding a pathway along which Jones does not freely vote for Gore and is not morally responsible for voting for Gore make it the case that Jones actually is morally responsible for voting for Gore? This is the "problem of alchemy" (Fischer 1994, p. 141).

Similarly, suppose one had a theory of knowledge according to which some individual S (the individual in question is always called "S"!) knows that p only if S can discriminate p from relevant alternatives. This is structurally analogous to the view that moral responsibility requires alternative-possibilities control. Whereas such a view is plausible, it would certainly be absurd to suppose that what transforms some case of lack of knowledge into a case of knowledge would be the existence of some alternative scenario in which the agent makes a mistake. How can adding a scenario in which S lacks knowledge (in this way) make it the case that S actually has knowledge (Fischer 1994, pp. 141–42)?

C. A Dilemma for the Proponent of Frankfurt-Type Examples

Regrettably (but not surprisingly) the above argument (that the flickers of freedom are insufficiently robust) has not been the final word on these matters! Indeed, a powerful challenge to the position I have sketched above (against the flicker theorist) has been presented by such philosophers as David Widerker (1995*a*, 1995*b*), Robert Kane (1996, esp. pp. 142–45), Carl Ginet (1996), and Keith Wyma (1997).[34] I will boil down the various versions of the argument into the following. It begins with a dilemma: the proponent of the Frankfurt-type examples is presupposing the truth of either causal determinism or indeterminism.

Let us start with the presupposition that causal determinism obtains. Now it does appear as if the relevant agent—Jones, in the example above—cannot choose or do otherwise (cannot choose at T_2 to vote for Bush or vote for Bush at T_3). This is because the "counterfactual intervener"—the liberal neurosurgeon Black—can know, given the prior

---

34. As far as I know, Kane was the first to articulate this strategy in reply to the Frankfurt examples; see Kane (1985), p. 51.

sign exhibited by Jones at T_1, that Jones will indeed choose to vote for Gore at T_2. If Jones were to choose at T_2 to vote for Bush, the prior sign would have had to have been different; thus, Jones cannot choose at T_2 to vote for Bush. But the problem is that the contention that Jones is morally responsible for choosing to vote for Gore and actually voting for Gore is put in doubt, given the assumption of causal determinism.

That is, if causal determinism is explicitly presupposed, it does not seem that someone could say that Jones is obviously morally responsible for his actual choice and action in a context in which the relationship between causal determinism and control (freedom) and moral responsibility are at issue. To do so would appear to beg the question against the incompatibilist.

Now suppose that indeterminism (of a certain relevant sort) obtains. Under this supposition it would not be dialectically inappropriate to claim that Jones is morally responsible for his actual choice at T_2 to vote for Gore and his vote for Gore at T_3. But now the contention that Jones cannot choose at T_2 to vote for Bush at T_3 is called into question. This is because there is no deterministic relationship between the prior sign exhibited by Jones at T_1 and Jones's subsequent choice at T_2. So, if we consider the time just prior to T_2, everything about the past can be just as it is consistently with Jones's choosing at T_2 to vote for Bush. Someone might think that if it takes some time for Jones to make the choice, Black can intervene to prevent the completion of the choice; but then Jones will still have the possibility of "beginning to make the choice," which is surely more robust than a mere flicker of freedom (say an involuntary twitch, blush, or neurological pattern). After all, beginning to make a choice is a voluntary undertaking (even if it is truncated through no fault of one's own).

The proponents of the Frankfurt-type examples contend that they are nonquestion-begging cases in which an agent is morally responsible for his choice and action and yet has no sufficiently robust alternative possibilities. But the counterargument of Widerker, Kane, Ginet, and Wyma appears to show that the examples in question are either not uncontroversial cases in which the agent is morally responsible for his choice and subsequent behavior or not cases in which the agent lacks the alternative possibilities. This is clearly an important argument, and it has been quite influential. Indeed, in a recent article Ted A. Warfield claims that the rejection of the Frankfurt-type examples (as cases in which an agent is morally responsible yet lacks alternative possibilities) is "increasingly common" (1996, esp. p. 221).

D. *A Reply on Behalf of the Proponent of the Frankfurt-Type Examples*

Despite this rising chorus I still remain convinced that the Frankfurt-type cases help to establish that it is very plausible that moral responsibility does not require alternative possibilities.

1. The Assumption of Causal Determinism

Begin with the first horn of the dilemma: the assumption that causal determinism obtains. I agree that one cannot now simply and precipitously conclude, from consideration of the examples, that the agent is morally responsible for his choice and behavior. But in any case this is not how I would have proceeded; I never have envisaged a simple one-step argument to the conclusion that (say) Jones is morally responsible for his choice and action. Rather, I employ the Frankfurt-type examples as the first (but obviously important) step of a slightly more complex argument to the conclusion that Jones is morally responsible for his choice and action (despite lacking alternative possibilities).

The argument goes as follows. First, one carefully considers the Frankfurt-type cases. On reflection, I believe that one should conclude that in these cases the lack of alternative possibilities does not in itself ground a claim that the agent is not morally responsible for his choice and action. In other words, I think that the examples make highly plausible the preliminary conclusion that if Jones is not morally responsible for his choice and action, this is not simply because he lacks alternative possibilities. After all, everything that has any causal (or any other kind of) influence on Jones would be exactly the same, if we "subtracted" Black entirely from the scene. And Jones's moral responsibility would seem to be supervenient on what has an influence or impact on him in some way.

So the relevant (preliminary) conclusion is, if Jones is not morally responsible for his choice and action, this is not simply because he lacks alternative possibilities. And it does not appear to beg the question to come to this conclusion, even if causal determinism obtains. The second step in the argument consists in asking whether causal determinism in itself and apart from ruling out alternative possibilities threatens moral responsibility. I have considered various possible reasons why someone might think that causal determinism does threaten moral responsibility in itself and apart from ruling out alternative possibilities, and I have come to the conclusion that it is not plausible to accept any of these reasons.[35] (I admit that I do not have any sort of decisive or "knockdown" argument for my conclusion here.)

2. The Assumption of Indeterminism

It seems to me that this two-stage argument is highly plausible and does not beg the question against the incompatibilist, even on the assumption of causal determinism. Thus I believe that the use of the prior-sign cases can be defended against the charge of begging the question. Let us now move to the second horn of the dilemma: the assumption of indetermin-

---

35. Fischer (1994), pp. 147–54. For further discussion of this issue, see Kane (1996), esp. pp. 40–43; and Mele (1996).

ism. Here I admit that the prior-sign cases may not be cases in which the agent does not have alternative possibilities. But I want to sketch three strategies for modifying the Frankfurt-type case to address this difficulty.[36]

i) *Hunt's approach.*—Recall that the original "Frankfurt-type" case was presented by John Locke in *An Essay concerning Human Understanding.* Locke's example is a case in which "a man be carried whilst fast asleep into a room where is a person he longs to see and speak with, and be there locked fast in, beyond his power to get out; he awakes and is glad to find himself in so desirable company, which he stays willingly in."[37] In Locke's example, the man stays in the room voluntarily, and it seems that he does so "freely" (although Locke himself would have used the term, 'voluntarily', rather than 'freely') and can be morally responsible for doing so, although he could not have left the room. Of course, the man does have various alternative possibilities (apart from special assumptions): he can choose to leave the room and try to leave the room, and so forth.

Frankfurt can be seen to be entering the debate at this point. Frankfurt seeks to construct examples in which even these sorts of alternative possibilities have been eliminated. To do this, Frankfurt employs the apparatus of a counterfactual intervener who can monitor the brain and intervene in it, should the agent be about to choose to do otherwise. In order to flesh out these examples—although Frankfurt did not explicitly do this—it is useful to posit a prior sign that can be read by the counterfactual intervener and guide him in his activity.[38] If the sign indicates that the agent is about to choose to do what the counterfactual intervener wants him to choose, the intervener does not intervene. If, contrary to fact, the agent were about to choose differently, the prior sign would inform the counterfactual intervener (and he would intervene).

A Frankfurt-type case which works as above is a prior-sign case. It is important to see that there can be another sort of Frankfurt-type case, which takes its cue more closely from Locke's example; I shall refer to such a case, developed by David Hunt, as a "blockage case."[39] Note that in Locke's example the door to the room is actually locked no matter whether the man is inclined to choose to stay in the room or not. Imagine, then, that although the actual neural processes in one's brain (one is here supposing that the mind supervenes on the brain) take place indeterministically, all other neural pathways are blocked.[40] This is a way of bringing the locked door—the blockage—into the head. Just as with

36. For yet another approach, see Fischer (1995); Widerker and Katzoff (1996); Hunt (1996); and Speak (1999).
37. Book 2, chap. 11, sec. 10.
38. This was Blumenfeld's innovation; see Blumenfeld (1971).
39. See Hunt (in press).
40. I borrow this example from Hunt (in press). He develops this—and related—examples further in his unpublished manuscript.

Locke's locked door, the pathways are actually blocked; in contrast to the structure of the prior-sign cases, the pathways' being blocked is not dependent on prior features of Jones. This, then, is a different way of solving precisely the problem Frankfurt sought to solve—one that more simply and naturally takes its cue from Locke. And, importantly, it does not appear to introduce alternative possibilities.

ii) *Mele and Robb's approach.*—Here is a second way of modifying the Frankfurt-type cases so that they (allegedly) "work" in a causally indeterministic context. Hunt's strategy involves "blockage" which is not sensitive to prior signs. The second strategy, developed by Alfred Mele and David Robb, involves two simultaneously operating sequences, one of which is indeterministic, the other of which is causally deterministic; the indeterministic sequence actually leads to the result in question, but the deterministic sequence (the operation of which is not sensitive to prior signs) would have issued in the same sort of result, if the indeterministic sequence had not. They develop their ingenious example as follows (changing our cast of characters slightly):

> At T_1, Black initiates a certain deterministic process P in Bob's brain with the intention of thereby causing Bob to decide at T_2 (an hour later, say) to steal Ann's car. The process, which is screened off from Bob's consciousness, will deterministically culminate in Bob's deciding at T_2 to steal Ann's car unless he decides on his own at T_2 to steal it or is incapable at T_2 of making a decision (because, e.g., he is dead by T_2). (Black is unaware that it is open to Bob to decide on his own at T_2 to steal the car; he is confident that P will cause Bob to decide as he wants Bob to decide.) The process is in no way sensitive to any "sign" of what Bob will decide. As it happens, at T_2 Bob decides on his own to steal the car, on the basis of his own indeterministic deliberation about whether to steal it, and his decision has no deterministic cause. But if he had not just then decided on his own to steal it, P would have deterministically issued, at T_2, in his deciding to steal it. Rest assured that P in no way influences the indeterministic decision-making process that actually issues in Bob's decision. (Mele and Robb, 1998, pp. 101–2)

The actual sequence in the Mele/Robb example is indeterministic, and yet the agent could not have done otherwise due to the unfolding of a deterministic causal sequence that preemptively overdetermines the actual decision.

iii) *Stump's approach.*—The third strategy for modifying the Frankfurt-type cases to accommodate indeterministic contexts is developed by Eleonore Stump.[41] Stump assumes that there is some sort of one-many correlation between a mental act or state and the firings of neurons in the brain:

41. Stump (1990, 1995, 1996*a*, 1999*a*, in press). Stump (in press) is a reply to Goetz (in press). For additional reflections, see Stump (1999*b*).

> When I suddenly recognize my daughter's face across a crowded room, that one mental act of recognition, which feels sudden, even instantaneous, to me, is correlated with many neural firings as information from the retina is sent through the optic nerve, relayed through the lateral geniculate nucleus of the thalamus, processed in various parts of the occipital cortex, which take account of figure, motion, orientation in space, and color, and then processed further in cortical association areas. Only when the whole sequence of neural firings is completed, do I have the mental act of recognizing my daughter. Whatever neural firings are correlated with an act of will or intellect, I take it that in this case, as in all others, the correlation between the mental act and the firing of the relevant neurons is a one-many relation. (In press, pp. 5–6)

On Stump's approach, it is crucial that if the firing of the whole neural sequence correlated with a mental act is not completed, the result is not some truncated or incomplete mental act (say, the beginning of a choice or decision). It is no mental act at all. She says:

> If the neural sequence correlated with my recognizing my daughter's face across a crowded room is interrupted at the level of the thalamus, say, then I will have no mental act having to do with seeing her. I won't, for example, think to myself, "For a moment there, I thought I saw my daughter, but now I'm not sure." I won't have a sensation of almost but not quite seeing her. I won't have a premonition that I was about to see her, and then I mysteriously just don't see her. I will simply have no mental act regarding recognition of her at all. (In press, p. 6)

Let us suppose now that a mental event is identical to a series of neural firings.[42] A particular mental event, say, a choice, can be assumed to be the result of an indeterministic process. Further, there can be a counterfactual intervener associated with the agent who could notice (in an alternative scenario) that a different neural sequence was beginning and could then interrupt it before it can be completed. If Black—the counterfactually intervening liberal neurosurgeon—did interrupt a neural sequence which was beginning to unfold (and which is such that, if it were completed, it would constitute—or correlate with—a decision to vote for Bush), Jones would not (according to Stump) have engaged in the mental act of beginning to make a decision. Jones would have no mental act, just as Stump would not have begun to recognize her daughter, if the sequence of neural firings beginning in her retina had been terminated in the thalamus (Stump, in press, p. 7).

Thus, in Stump's version of the Frankfurt-type cases, the agent's choice is not causally determined, and it is also true that the agent can-

---

42. This supposition is just for simplicity's sake; Stump's view is compatible with other stories as to the precise relationship between mental states and brain events.

not have chosen (or behaved) differently from how he actually chooses (and behaves). And yet it seems entirely plausible that the agent be morally responsible for his choice and behavior in these cases.

Despite the force and influence of the argument (presented by Widerker, Kane, Ginet, and Wyma) against the contention that in the Frankfurt-type cases the agent is morally responsible although he has no alternative possibilities, there is an attractive strategy of response. Even if causal determinism is true, it does not appear to be question-begging to use the cases as part of a two-stage argument (rather than an argument that simply assumes that the relevant agents are morally responsible in the cases). And if causal determinism is false (in certain ways), it still seems to be possible to construct versions of the Frankfurt-type cases in which it is plausible to say that the agent is morally responsible and yet lacks alternative possibilities.

E. A Reply on Behalf of the Flicker Theorist: Escapability of Authorship/Responsibility

Perhaps it will not shock the reader to learn that this is not the end of the story. The critic of the Frankfurt-type examples has (at least) one more card to play. Consider, for example, the following remarks of Michael McKenna: "Here I believe that Fischer has not fully addressed what motivates the advocate of [the alternative possibilities control requirement]. . . . What intuitively drives [the proponent of this requirement] is the kind of control needed in order for us to avoid being the author of a *particular* act and thus avoid being responsible for the production of *that* particular action. . . . It is a matter of holding people accountable for what they do only if they can avoid any blame or punishment that might fall upon them for performing those very particular actions which they do perform" (1997, pp. 73–74; emphasis added). McKenna elaborates as follows:

> The issue . . . here is whether the will . . . places *my* stamp upon the world, and whether *it is up to me . . . to have that particular stamp or some other as my mark upon the world.* In the Frankfurt-type cases the alternatives are, either doing what one does of one's own intention, or being coerced into performing the same kind of action against one's will. These alternatives do seem to be quite impoverished; however, they mean all the difference between one's doing something of one's own will, and one's not doing that kind of thing of one's own will. . . . What more fundamental kind of control can there be here other than the control for one to either have a particular will or not have it? (1997, pp. 74–75; emphasis added)

McKenna is claiming that even in the Frankfurt-type cases, the relevant agent has a significant and robust power: the power either to be the author of his action or not and thus the power to be morally responsible

for his action or not. A similar point is made in an interesting recent article by Keith Wyma (1997). Wyma begins with an example which suggests that many of us experienced something like a Frankfurt-type example as we were growing up:

> When I was four years old and learning to ride a bicycle, I reached a point where my father decided I no longer needed training wheels. But he still worried that I might fall. So on my first attempt "without a net," he ran alongside as I pedaled. His arms encircled without touching me, his hands resting lightly upon me, but not holding me upright. I rode straight ahead. My father did not push or guide me, but if I had faltered or veered suddenly to the side, he would have tightened his grip, keeping me vertical and on track. After finally braking to a stop, I was jubilant but somewhat hesitant over whether I should be. I wondered, had I really ridden my bike on my own? . . . Was the triumph of riding straight down the street mine or not? (1997, p. 57)

Wyma goes on to argue for an intuition very similar to McKenna's. On Wyma's view, moral responsibility requires a certain kind of "leeway." And this leeway is specified by what Wyma calls the "principle of possibly passing the buck" (PPPB): "A person is morally responsible for something she has done, *A,* only if she has failed to do something she could have done, *B,* such that doing *B* would have rendered her morally nonresponsible for *A*" (1997, p. 59). Of course, in a Frankfurt-type case the relevant agent would not be morally responsible in the alternative sequence; Jones would not be morally responsible for voting for Gore in the circumstance in which Black's device were triggered. Thus Wyma has apparently identified a significant sort of leeway, even in the Frankfurt-type examples. At the end of his paper, Wyma returns to the analogy with which he started, saying, "I believe the bike riding triumph *was* mine, because even though I could not have fallen or crashed while my father hovered protectively over me, I could still have faltered enough that he would have had to steady me; and because I had leeway to falter but did not do so, the success of riding was truly mine. *PPPB* vindicates a similar kind of leeway as being necessary for ascriptions of moral responsibility" (1997, p. 68).

Additionally, Michael Otsuka has recently defended a principle similar to Wyma's principle of possibly passing the buck. Otsuka calls his principle, the "Principle of Avoidable Blame": "One is blameworthy for performing an act of a given type only if one could instead have behaved in a manner for which one would have been entirely blameless." [43]

Thus, all three defenders of the alternative-possibilities control requirement seem to be pointing to the same sort of alternative possibility

---

43. Otsuka (1998), esp. p. 688. Otsuka qualifies the principle to apply to cases in which it is not the case that everything one is capable of doing at a given point in time is blameworthy because of some previous choice for which one is to blame.

which they claim is present quite generally, and hence in the Frankfurt-type examples. This is the freedom to "pass the buck" or "escape" or "avoid" moral responsibility.[44] And it seems that this freedom is present in all of the modifications of the Frankfurt-type examples presented above. One might say that these theorists are seeking—perhaps with some success—to fan the flickers of freedom.[45]

F. A Reply (Again) on Behalf of the Proponent of the Frankfurt-Type Examples

But before we mistake these breezes for Santa Ana winds, consider the following replies. First, it may be that David Hunt's approach shows that there can be cases in which an agent is morally responsible for his choice and behavior, and yet he lacks even the sort of alternative possibility now under consideration: the possibility of escaping authorship or responsibility. Recall that Hunt envisages a case in which the neural events resulting in the relevant choice are indeterministic, and yet all other neural pathways in the brain are "blocked" (as in Locke's "locked-door" example). The question could now be put as follows: Does the agent have access to a scenario in which his neural path makes contact with or "bumps up against" the blockage? If so, it would seem that the alternative possibility in question does exist, after all, because if the neural path bumps up against the blockage, then presumably the agent is no longer the author of the subsequent act (and is not morally responsible for it).

But how exactly can the agent (or his neural events) bump up against the blockage? It would seem that in order to have access to the blockage, there would have to be an intermediate set of neural events, different from the actual neural events, that is, as it were, a "bridge" between the actual neural process and the blockage. (In Locke's example, the agent would have to walk over to the door and try to open it.) But even these intermediate events are presumed to be blocked in Hunt's example. So it may seem that Hunt has provided, indeed, an example of the required sort, that is, one in which the agent is morally responsible and yet does not have any alternative possibilities (even the possibility of avoiding authorship and moral responsibility). But the example is difficult to imagine (and thus properly to evaluate). If causal indeterminism obtains in the actual neural pathway, how exactly can it be the case that the agent does not have access to events consisting in bumping up against the barriers?[46]

---

44. It is interesting to note that Wyma seems to be adopting Watson's "genuine imputability" notion of moral responsibility, whereas Otsuka seems to be adopting Watson's "accountability" notion. McKenna may be considering moral responsibility in both senses.

45. Thanks to Dan Speak for this phrase.

46. To help with this point, David Hunt has suggested the following analogy in personal correspondence:

> A vertical pipe fills with water and freezes; the sun thaws the pipe and a plug of ice moves down the pipe. Alternative directions are blocked by the pipe, but the ice

Hunt has also suggested that the context of God's foreknowledge of future events is relevantly similar to Frankfurt-type examples. Let us suppose that God exists within the same time framework as humans, is essentially omniscient, and can know future contingent truths. Let us further assume that causal indeterminism obtains. (Of course, each of these assumptions is contentious, as is their combination.) I believe that it follows from the conjunction of these assumptions (suitably interpreted) that human agents cannot choose or do otherwise, and yet (given certain assumptions about God) God's knowledge plays absolutely no role in human choices and actions. Just as with the counterfactual intervener in a Frankfurt-type case, one could "subtract" God from the situation and everything that has a causal impact on the agent's choices and behavior would be exactly the same. If all the above is correct, then the context of God's foreknowledge would seem to be one in which an agent could be held morally responsible for his choice and behavior and yet have no alternative possibilities (even the possibility of avoiding authorship and responsibility).[47]

Second, I believe that problems similar to the problems with the earlier defenses of the alternative-possibilities control requirement also plague the new approaches. Recall that the problem with saying that it is the possibility of exhibiting a different prior sign or indicator of future decision (and action) that grounds moral responsibility is that the envisaged possibility is too exiguous and flimsy. The displaying of such a sign would not even be voluntary behavior. How could moral responsibility rest on such a delicate foundation?

Now it might be thought that the possibility of avoiding authorship or the possibility of avoiding moral responsibility would be a more substantial basis for moral responsibility. But I believe there are similar prob-

---

never bumps up against the blockage—at least in the way that a marble rattling down the pipe would do. Let's develop the example so that it gives us the right parallel. The ice plug has to be the will, not a body with a will "inside" it. Since we are supposing that the will is causally indeterministic, imagine that the ice plug is moving vertically under the influence of gravity, but that at one-second intervals it might (or might not) shift indeterministically on the horizontal plane (without ceasing to move vertically). We observe its movements over a 10-second period. Suppose its position during each of these one-second intervals is as follows (the intervals are numbered from 1 to 10): 1,A; 2,B; 3,B; 4,C; 5,D; 6,E; 7,E; 8,E; 9,E; 10,F. Because the ice plug's position is brought about indeterministically, it is natural to think that at time 7 it could just as well have ended up in, e.g., position G, rather than position E. But suppose that at time 6 the ice plug happened to align itself perfectly with the mouth of a perfectly translucent and perfectly frictionless pipe, which it then entered as it continued to move vertically under the force of gravity. (It exits the pipe at time 10.) The ice plug's position at times 7–9, we are supposing, is actually brought about indeterministically, in the same manner as its position was brought about at times 1–6. Nothing changes in this respect. But the presence of the pipe ensures that alternative positions are unavailable at times 7–9.

47. For an extended discussion, see Hunt (1999).

lems here.[48] Note that in the alternative sequence in a Frankfurt-type case the agent would indeed be avoiding (say) moral responsibility, but he would be doing so "accidentally." The agent would not be voluntarily avoiding responsibility. The suggestion that avoiding responsibility is a sufficiently robust basis for moral responsibility may get some of its plausibility from the fact that in a typical context in which we would say that someone has avoided (say) blameworthiness, it would be in virtue of some voluntary action. Typically, the relevant facts about the various paths available to the agent would be accessible to him, and he would voluntarily choose a right action (rather than a morally objectionable one). Here we would say that the agent avoided blameworthiness; but this is a very different sort of context from the Frankfurt-type cases. In the Frankfurt-type cases, the agent does not choose to be morally responsible rather than not—these issues play no role in his deliberations. And in the alternative scenario in a Frankfurt-type case, the agent does not choose to escape responsibility or voluntarily choose anything which implies his escaping responsibility.

To bring this point out a bit more clearly, note that in the alternative scenario in a Frankfurt-type case the agent does not deliberate about whether or not to embrace moral responsibility. Issues about whether or not to be morally responsible play no explicit role in his deliberations. Further, they play no "implicit" role either. They might play an implicit role in the sort of context discussed above in which an agent has internalized certain norms on the basis of which he chooses to do what he takes to be the right action. If he successfully avoids blameworthiness here, it is partly in virtue of his having internalized norms the relevant community shares. Given these norms, the agent can reasonably expect to escape blame, if he chooses as he does. But in the alternative scenarios in the Frankfurt-type cases, issues about moral responsibility obviously do not play an implicit role of this sort.

To the extent that issues pertaining to moral responsibility play neither an explicit nor an implicit role, I say that moral responsibility is not "internally related" to the agent's behavior in the alternative sequence of a Frankfurt-type case. And my point is that it is very plausible that moral responsibility must be so related to the agent's behavior in order for the alternative possibility in question to be sufficiently robust to ground ascriptions of moral responsibility.

Of course, I do not accept the alternative-possibilities control model of moral responsibility. But my contention is that, if you do buy into this traditional picture, then you should also accept that the alternative possibilities must be of a certain sort—they must be sufficiently robust. This same point has been highlighted by a philosopher with a very different orientation from mine: Robert Kane (1985, esp. p. 60; 1996, esp.

48. Here again I borrow from Fischer (in press).

pp. 107–15). (Kane is a libertarian who believes that alternative possibilities are required for moral responsibility.) Kane emphasizes what he calls the "dual" or "plural" voluntariness (and responsibility) conditions on moral responsibility: the relevant alternative possibilities—that is, alternative possibilities sufficiently robust to ground moral responsibility—must themselves involve voluntary behavior (for which the agent is morally responsible). On Kane's picture, it is not enough that an agent have just any sort of alternative possibility; it must be an alternative in which the agent acts voluntarily and is morally responsible. Similarly, I would contend that the relevant alternative possibilities must contain voluntary, responsible behavior in which moral responsibility is internally related to the agent's behavior. My suggestion, then, is that the new defenses of the alternative-possibilities control requirement (presented by McKenna, Otsuka, and Wyma) fall prey to the same sort of problem that afflicted earlier such defenses: the alternatives they postulate are not sufficiently robust.

In my early article, "Responsibility and Control," I argued that the critic of the Frankfurt-type case mixes up "possibility" and "ability" in a certain way. That is, I pointed out that even if another event (or set of events) occurs in the alternative sequence of a Frankfurt-type case, it does not follow that the agent has the ability (in the relevant sense) to bring this alternative event (or set of events) about. I believe that the recent defenses of alternative-possibilities control simply reinscribe the same general problem. The lack of "internal relatedness" of moral responsibility to the events in the alternative sequence points to the fact that the agent lacks the relevant sort of ability, even if there exists the possibility of something different happening. Even if there exists the possibility that the agent not be the author of his action (or avoid moral responsibility), it does not follow that the agent has the ability (in the relevant sense) to avoid authorship (or responsibility). It is a simple point that has played a crucial role in discussions of indeterministic conceptions of control and moral responsibility: the mere possibility of a different event occurring does not entail that the agent has the ability to do otherwise. (The mere possibility of a different event occurring is not sufficient for "agent control," as the rollback argument shows.) The point applies equally in the context of the Frankfurt-type examples.

Return to Wyma's striking claim about his early bike-riding experience: "I believe the bike riding triumph *was* mine, because even though I could not have fallen or crashed while my father hovered protectively over me, I could still have faltered enough that he would have had to steady me; and because I had leeway to falter but did not do so, the success of riding was truly mine" (1997, p. 68). Whereas we could quibble endlessly about details of these sorts of examples, it seems to me that the intuitive point is quite clear: it is not the possibility of faltering slightly

that makes the young Wyma's bike riding triumph truly his. This has to do not with whether he could have faltered slightly but with how he rode the bike—how he moved the pedals, balanced, and so forth, and by what sort of causal process this all took place.

In Frankfurt-type cases, an agent is morally responsible for his action, although he lacks the relevant kinds of alternative possibilities. He cannot, then, make a relevant difference to the world; he does not (in the appropriate way) select one path for the world to take among various genuinely open paths. But the agent is, nevertheless, fully and robustly morally responsible for what he does.[49]

G. Some Putative Implications of the Frankfurt-Type Cases

There has (obviously) been considerable ink spilled over the Frankfurt-type examples. Recently some philosophers have explored some previously unnoticed (or insufficiently noticed) aspects of the examples. The Frankfurt-type examples purport to be contexts in which an agent can be morally responsible even though he lacks alternative possibilities. But consider the following argument, which has been presented and discussed recently by David Widerker (1991), David Copp (1997), and Ishtiyaque Haji (1993).[50] Suppose someone does something which is intuitively "bad," such as lying just to bolster his reputation. If this act is blameworthy, then it must be wrong. And if it is wrong, it must be the case that the agent should have done something else instead (where this could include simply refraining from doing anything). But "ought implies can," so if the agent should have done something else instead, then he must have been able to do something else. Thus, if the agent had no alternative possibilities (and thus could not have done anything else), then his act of lying cannot be considered blameworthy.

The above sort of argument threatens the idea that an agent can be genuinely blameworthy in a context in which he has no alternative possibilities. Insofar as an account of moral responsibility will certainly need to accommodate agents being blameworthy on some occasions, the

49. I argued above that the intuitive picture behind the alternative-possibilities control requirement on moral responsibility is the idea that one must make a difference to the world through one's behavior by selecting one from among various genuinely open pathways into the future. Since the Frankfurt-type cases show that moral responsibility need not involve an agent's making such a difference to the world, a different intuitive picture of moral responsibility is required. I offer some tentative work toward such a picture in Fischer (in press). Here I argue that although the agent need not make a difference in order to be morally responsible, he does make a statement of a certain sort. Thus, I argue for a certain sort of self-expression model of moral responsibility. I believe that the "value" of acting freely and thus acting so as to be morally responsible consists in making a certain kind of statement, not in making a difference.

50. There is an instructive, extended discussion of these issues in Haji (1998), esp. pp. 42–64 and 151–67.

argument calls into doubt whether an adequate "actual-sequence" account of moral responsibility can be given.[51]

The argument is disturbing, and it is worthy of more careful attention than I can give it here. I can here only briefly suggest various ways of responding to the argument. One might deny the maxim that "ought implies can." Various philosophers have rejected this maxim on grounds quite independent of considerations pertinent to the Frankfurt-type cases. Typically these philosophers have been motivated to give up this maxim in light of reflection on the logic of moral dilemmas.[52] Another approach is suggested by Haji, who rejects the contention that if an act is blameworthy, then it is "objectively wrong." That is, the argument presupposes that blameworthiness is connected to objective wrongness, whereas Haji believes that blameworthiness is linked only with subjective wrongness. So, on Haji's view, if an agent is blameworthy for performing an action, it need not be the case that the action was wrong, only that the agent believed it to be wrong (and nevertheless did it).[53]

My own inclination here is to reject the "ought implies can" maxim. The maxim says that if an agent ought to do *X,* then he can do *X.* But why exactly should one accept this maxim? That is, what justification could be offered for this maxim? It is most natural, I think, to say that the maxim is valid because if it were not, then there could be cases in which an agent ought to do *X* but in fact cannot do *X* (and never could do *X*). Thus, given the connection between its being the case that an agent ought to do *X* and the agent's being blameworthy for not doing *X,* there could be cases in which an agent is blameworthy for not *X*-ing and yet he cannot *X.* And this seems unfair.

But I believe that there are Frankfurt-type omissions cases that are relevantly similar to Frankfurt-type cases with respect to actions. That is, there are cases in which an agent is morally responsible for not *X*-ing although he cannot in fact *X.*[54] Some of these are cases in which an agent is blameworthy for not *X*-ing and yet he cannot *X.* In fact, I believe that anyone who accepts the Frankfurt-type action cases must accept that there are such omissions examples. Thus, it is precisely the basic intuitions elicited by the Frankfurt-type cases which show that the most natural justification of the "ought implies can" maxim is faulty. It is therefore

51. Haji has employed a similar argument to call into question whether morality itself could exist in a world without alternative possibilities; see ibid., pp. 42–54.

52. For a thorough discussion, see Sinnott-Armstrong (1988). Sinnott-Armstrong believes that the "ought implies can" maxim is not an entailment but functions like a conversational implicature.

53. There is a critical discussion of this view in Copp (1997). For an alternative way of challenging the Widerker/Copp argument, see Yaffe (1999*b*).

54. I and my coauthor argue for this claim in Fischer and Ravizza (1998), esp. pp. 123–50.

not ad hoc for anyone who accepts the standard interpretation of the Frankfurt-type cases to reject the "ought implies can" maxim.[55]

IV. ACTUAL-SEQUENCE ACCOUNTS OF MORAL RESPONSIBILITY

A. *Identification Approaches*

In the previous section I have defended the idea that moral responsibility need not require alternative-possibilities control; that is, I have defended the notion that the appropriate account of moral responsibility will be an actual-sequence account.[56] It is instructive to distinguish two sorts of actual-sequence models of moral responsibility. The first such approach is the "identification" approach to moral responsibility.

i. *Frankfurt's hierarchical account.*—Various philosophers have offered what might be called identification accounts of moral responsibility. One of the most salient such accounts has been offered (in different forms over the years) by Harry Frankfurt (1971, 1987, 1992). Frankfurt's approach is "hierarchical" in the sense that it employs the apparatus of higher-order preferences: preferences that have as their objects other preferences (rather than courses of action or states of affairs apart from motivational structures).

In his early work, Frankfurt seemed to suggest that having a "second-order volition" to act in accordance with the first-order desire that actually moves one to act ("the will") is sufficient for identification, acting freely, and moral responsibility (at least with respect to the "freedom-relevant" component of moral responsibility). But this account faces various problems. For example, it might be that the agent has conflicting second-order volitions. Or it might be that the agent doesn't care about his second-order volition.[57]

To address these sorts of problems, Frankfurt has developed two separate strategies of response (in his later articles).[58] The first strategy involves adding "wholehearted decision" to the hierarchical "mesh" noted above (the conformity of the second-order volition to the will).

55. I thank Mark Ravizza for helping me to see this point. Note that the Widerker/ Copp argument is a challenge for any account of moral responsibility according to which responsibility does not require alternative possibilities. For example, it is a challenge for a theory such as R. Jay Wallace's, as well as those that are motivated by Frankfurt-type examples. It is interesting that (as I pointed out in the text) consideration of the Frankfurt-type cases provides a powerful way of rejecting the Widerker/Copp argument; it is not clear that one who wishes to eschew such examples has a similarly potent response.

56. By an "actual-sequence" account I simply mean an account of moral responsibility which does not require the availability of alternative possibilities (of the sort corresponding to Austin's "all-in" sense of 'can'). I do not mean that such an account cannot make use of alternative scenarios (or other possible worlds), perhaps as a means of fixing the modal characteristics of the actual sequence.

57. For this sort of criticism, see Watson (1975).

58. Here I have been helped by Bratman (1996, 1999*a*).

Here it is required (roughly) that the relevant higher-order volition be unopposed by a volition of the same order, and that the agent judge that no further consultation with even higher-order preferences would lead to a reversal of his ordering of preferences.

Another strategy appears to reject the requirement of wholehearted decision in favor of what might be called a "satisfaction" condition: "identification is constituted neatly by an endorsing higher-order desire with which the person is satisfied" (Frankfurt 1992, p. 14). Frankfurt elaborates, "being genuinely satisfied is not a matter, then, of choosing to leave things as they are or of making some judgment or decision concerning the desirability of change. It is a matter of simply *having no interest* in making changes. What it requires is that psychic elements of certain kinds *do not occur*" (1992, pp. 13–14).

But I believe that there are significant problems with both approaches, insofar as they purport to offer sufficient conditions for the freedom-relevant component of moral responsibility. As regards Frankfurt's "wholehearted decision," it seems that this sort of mental state could be induced by a "demonic neurologist" via a responsibility-undermining process (just as much as a mesh between the higher-order volition and the will could be). Thus, a wholehearted decision, in combination with the hierarchical mesh condition, does not yet provide a sufficient condition for moral responsibility.

Further, Bratman has pointed out that Frankfurt's account of "satisfaction" seems to render the notion open to the following sort of worry. The mere absence of an interest in change, even if (as Frankfurt requires) this is grounded in some sort of reflection, may be due to "ennui, depression, or exhaustion, or enervation, or the like" (Bratman 1996, p. 7). Thus, the notion of satisfaction, understood as Frankfurt understands it, seems too weak to capture the idea of identification that is relevant to acting freely and to moral responsibility.

ii. *Other identification models.*—There are various ways of seeking to embellish Frankfurt's minimalist hierarchical apparatus with elements that may get one closer to an adequate account of moral responsibility. One general approach requires that the hierarchical mesh (the conformity between the relevant higher-order preference and the will) be based on value judgments of an appropriate kind. One might call this the "value-added" approach.[59]

Another approach is suggested by Michael Bratman, who recently has combined the minimal hierarchical elements with more robust ingredients from his "planning approach" to agency.[60] In a series of papers

---

59. Thanks to Andrew Eshleman for this term. See Taylor (1976); and Stump (1988, 1996*b*, 1997). Also, see Eshleman (1998).

60. Bratman develops the planning approach in (1987, 1999*b*).

(1996, 1997, 1999*b*) Bratman has been developing a theory that seeks to address some of the challenges to the identification approach (especially as developed by Frankfurt). Bratman puts the problem in terms of identifying which desires are more centrally part of the acting self, and he says (by way of summarizing the overall approach), "the proposal is that desires that are part of the constitution of the acting self are ones endorsed by relevant decisions, intentions, and policies; and it is some such decision, intention, or policy in favor of a desire that constitutes the justificatory status, for the agent, of that desire. Finally, we say in general what such decisions and intentions are by appeal to the planning theory. In this way we combine the identification strategy with a modest theory of the will grounded in the planning theory" (1999*b*, p. 23). Bratman believes that the more specific account he ultimately develops gives an analysis of the "aretaic" or "genuine imputability" aspect of moral responsibility (in Gary Watson's terms). One would need to combine this analysis with a separate account of the "accountability" aspect of moral responsibility to get a full explanation of moral responsibility.[61]

B. Reasons-Responsiveness Accounts

A second sort of "actual-sequence" model of moral responsibility connects moral responsibility with "responsiveness to reasons." Again, there are various different versions of this kind of strategy.[62] R. Jay Wallace contends that what makes it fair to hold an agent to moral obligations (in the absence of an excuse in the particular context in question) is his possession of the general powers of reflective self-control (1994). Roughly speaking, this capacity involves the ability to grasp moral reasons and control one's behavior in light of them.

My coauthor, Mark Ravizza, and I also develop a reasons-responsiveness approach to the freedom-relevant component of moral responsibility. On our approach one distinguishes between the way the actual

61. Although Bratman does not explicitly address the issues of whether moral responsibility requires alternative possibilities or whether it is compatible with causal determinism, his account of the aretaic component of responsibility appears not to require alternative possibilities or to be inconsistent with causal determinism.

62. As I pointed out above, Susan Wolf develops what has been called a "normative" approach to freedom and moral responsibility. On her view, a certain sort of responsiveness to reasons (or "normative competence") is required for moral responsibility. That is, Wolf's "reason view" has it that "the freedom necessary for responsibility consists in the ability (or freedom) to do the right thing for the right reason . . . to choose and to act in accordance with the True and the Good" (Wolf 1990, p. 94). It may appear as if this is not an actual-sequence approach to moral responsibility, since it seems that Wolf contends that moral responsibility for wrong acts requires alternative possibilities. But it is not clear that the alternative possibilities in question (the "abilities" Wolf has in mind) are the "genuine" possibilities associated with Austin's "all-in" sense of 'can' and the incompatibilist's argument that causal determinism rules out alternative possibilities.

sequence unfolds (the actual-sequence mechanism) and the processes that take place in various alternative scenarios.[63] This is crucial to addressing the issues posed by the Frankfurt-type cases. One then requires that the actual-sequence mechanism be suitably sensitive to reasons (some of which must be moral reasons).[64]

R. Jay Wallace and Fischer and Ravizza link moral responsibility with control. But we are concerned to distinguish two kinds of control. Moral responsibility, on our approaches, does not require alternative-possibilities control. On our approaches, it does involve a sort of control that is analyzed in terms of responsiveness to reasons. On Wallace's model, the control in question is possessed by an agent with certain general powers of reflective self-control and is manifested by such an agent in a context in which he has no excuse for his behavior. On the Fischer and Ravizza model, the manifestation of this sort of control consists in the agent's actually acting on his own, suitably reasons-responsive mechanism.[65] Thus if causal determinism threatens alternative-possibilities control, it does not thereby threaten the approaches of Wallace and Fischer and Ravizza.

63. For the sake of convenience we employ the term, 'mechanism', but we mean by (say) 'actual-sequence mechanism' the way the actual sequence unfolds or the kind of process that takes place in the actual sequence.

64. For initial sketches of the approach, see Fischer (1987, 1994); for a more detailed development, see Fischer and Ravizza (1998). Although we have considerable sympathy with Wallace's overall project and are in general agreement with Wallace's approach, we do not (in contrast to Wallace) simply require that the agent possess the relevant general capacity (without necessarily having "access" to it). (For critical discussions of this aspect of Wallace's theory, see Fischer [1996]; and Clarke [1997].) On our approach, an agent's actual-sequence mechanism must be suitably reasons-responsive; thus, in behaving as he does, the agent exercises the general capacity in question. The Fischer and Ravizza approach can then be put "in between" the incompatibilist's claim that moral responsibility requires the genuine availability of alternative possibilities and Wallace's view that it simply requires the possession of the relevant general capacity to recognize and respond to reasons (given that there is no excuse in the particular context in question). The Fischer and Ravizza approach does not require genuine alternative possibilities, but it does require that the agent have access to the general capacity in question: it requires that the agent exercise this capacity (by acting on a suitably reasons-responsive mechanism). But note that although our approach entails access to the general capacity in question, it does not thereby require access to the possible worlds by reference to which the general capacity is defined; put in other words, whereas the Fischer/Ravizza approach requires that the actual sequence display certain modal characteristics and that the relevant agent have access to the mechanism with those characteristics, it does not require that the agent have access to the possible worlds in terms of which the modal characteristics are defined: see Fischer and Ravizza (1998), pp. 51–54, 143–44, n. 31.

65. Presumably a demonic neurosurgeon could manipulatively implant a reasons-responsive mechanism (just as he could induce a "mesh" between a higher-order motivational state and the agent's will); that is why we require that the mechanism be the agent's own. For an account of mechanism ownership, see ibid. esp. pp. 207–39.

V. THE TAO OF SEMICOMPATIBILISM

I want to end by sketching what I take to be a very powerful motivation for embracing "semicompatibilism"—the doctrine that causal determinism is consistent with moral responsibility, even if causal determinism rules out alternative possibilities. I believe that we—you and I and most adult human beings—are morally responsible (at least much of the time) for our behavior. Further, I do not think that this very important and basic belief should be "held hostage" to esoteric scientific doctrines. For example, if I were to wake up tomorrow and read in the *Los Angeles Times* that scientists have decisively proved that causal determinism is true, I would not have any inclination to stop thinking of myself, my family and friends, and human beings in general as morally responsible. The precise form of the equations that describe the universe, and whether or not they are or correspond to universal generalizations, are not the sorts of thing that should be relevant to our most basic views of ourselves (as morally responsible agents and thus apt targets of the reactive attitudes).

I cannot imagine, for example, reading that some Cal Tech scientists have finally established that the equations that describe the universe are indeed deterministic (rather than, say, "almost deterministic") and then concluding that I cannot have certain distinctive attitudes—such as love, hatred, respect, gratitude, and so forth—toward my family and friends (and in general other human beings). I certainly cannot imagine simply concluding that there are no deep differences (relating to the reactive attitudes) between human beings and other animals (and inanimate objects). Our reactive attitudes should not be held hostage to an esoteric scientific discovery of the kind in question. That is, the reactive attitudes, and our views of ourselves as morally responsible agents, should be resilient in a certain sense.

This resiliency idea is a major motivation for my acceptance of semicompatibilism. It is part of the background against which I evaluate the complicated debates pertaining to the Frankfurt-type cases, and it makes me more inclined to conclude that such cases do indeed establish that alternative possibilities are not required for moral responsibility. It also influences my evaluation of the question of whether causal determinism in itself and apart from considerations pertinent to alternative possibilities rules out moral responsibility.

What may be surprising is that my acceptance of the resiliency idea and semicompatibilism makes me close allies with the indeterminists I described in Section II (such as Ginet, Kane, O'Connor, and Clarke). Of course, we do not agree about the requirement of alternative possibilities, and I do not accept that moral responsibility entails that causal determinism must be false.[66] But there can be disagreements among friends! And

---

66. Randolph Clarke is inclined in fact to agree that moral responsibility does not require alternative possibilities, but he does insist that it requires the falsity of causal deter-

we are friends to the extent that we—both camps—accept that the falsity of causal determinism in itself would not rule out control and moral responsibility. For a proponent of the resiliency idea must surely hold that our moral responsibility cannot be held hostage to the discovery of the falsity of causal determinism (any more than to the discovery of the truth of causal determinism). I would certainly feel no inclination to give up my view of human beings as morally responsible if a consortium of scientists decisively established that causal determinism is false (in a way that falls short of randomness). Although I do not have any distinctive perspective on the challenge posed by the "rollback argument" to the indeterminist to show how an agent's actions in an indeterministic world can genuinely be an "outflowing" of the agent, I am committed in principle to the notion that this challenge is not insuperable.

In contrast to the libertarian, I need not claim that causal determinism must be false in order for us to be morally responsible.[67] But with the libertarian, I do believe that the falsity of causal determinism would not in itself show that we lack the kind of control associated with moral responsibility. Due to the motivating engine of resiliency, I am then a kind of "supercompatibilistic semicompatibilist."

REFERENCES

Adams, Robert. 1985. Involuntary Sins. *Philosophical Review* 94:3–31.
Arenella, Peter. 1990. Character, Choice, and Moral Agency: The Relevance of Character to Our Moral Culpability Judgments. *Social Philosophy and Policy* 7:59–83.

minism. He, like Stump and David Hunt, is a "hyper-incompatibilist"; he believes that causal determinism rules out moral responsibility even though moral responsibility does not require alternative possibilities.

67. I am thus not subject to a weird sort of dialectic of flip-flopping to which Peter Van Inwagen is vulnerable. (See Fischer and Ravizza 1998, esp. pp. 253–54.) Van Inwagen is very confident that we are morally responsible. And yet he is also very confident that certain metaphysical principles are valid and that they entail the incompatibility of causal determinism and moral responsibility. So he concludes that causal determinism is false. But Van Inwagen also says that in the unlikely scenario in which he were convinced that causal determinism were true, he would not give up his view that we are morally responsible; rather, he would jettison at least one of the relevant metaphysical principles. But it seems bizarre that the tenability of such a principle should depend on the empirical thesis of causal determinism. In contrast to a libertarian such as Van Inwagen, I would not need to reconsider my relevant metaphysical principles (such as the "fixity of the past" and the "fixity of the natural laws") under the supposition that causal determinism were discovered to be true. I am not vulnerable to this sort of metaphysical flip-flopping—or to the equally unappealing possibility of having to give up my view of human beings as morally responsible agents. Now, of course, I do not deny that my general desire to protect the resiliency of the reactive attitudes may marginally influence my evaluation of various philosophical claims, but this is crucially different from supposing that the truth of an empirical doctrine such as causal determinism should bear on the acceptability of metaphysical doctrines such as the "fixity of the past," the "fixity of the laws," or certain modal principles having to do with the "transfer of powerlessness" or the "transfer of nonresponsibility."

Aristotle. 1985. *Nicomachean Ethics.* Translated by Terence Irwin. Indianapolis: Hackett Publishing Co.

Audi, Robert. 1986. Acting for Reasons. *Philosophical Review* 95:511–46.

Audi, Robert. 1991. Responsible Action and Virtuous Character. *Ethics* 101:304–21.

Audi, Robert. 1993. *Action, Intention, and Reason.* Ithaca, N.Y.: Cornell University Press.

Beaty, Michael D., ed. 1990. *Christian Theism and the Problems of Philosophy.* Library of Religious Philosophy, vol. 5. Notre Dame, Ind.: University of Notre Dame Press.

Bennett, Jonathan. 1980. Accountability. In *Philosophical Subjects: Essays Presented to P. F. Strawson,* ed. Zak van Straaten, pp. 14–47. Oxford: Clarendon Press.

Benson, Paul. 1987. Freedom and Value. *Journal of Philosophy* 84:465–86.

Benson, Paul. 1990. The Moral Importance of Free Action. *Southern Journal of Philosophy* 28:1–18.

Bernstein, Mark. 1992. *Fatalism.* Lincoln: University of Nebraska Press.

Berofsky, Bernard. 1987. *Freedom from Necessity: The Metaphysical Basis of Responsibility.* New York: Routledge & Kegan Paul.

Berofsky, Bernard. 1995. *Liberation from Self: A Theory of Personal Autonomy.* Cambridge: Cambridge University Press.

Bishop, John. 1989. *Natural Agency: An Essay on the Causal Theory of Action.* Cambridge: Cambridge University Press.

Blumenfeld, David. 1971. The Principle of Alternate Possibilities. *Journal of Philosophy* 67:339–44.

Bogdan, R. J., ed. 1982. *Roderick M. Chisholm.* Dordrecht: D. Reidel Publishing Co.

Bok, Hillary. 1998. *Freedom and Responsibility.* Princeton, N.J.: Princeton University Press.

Bradie, Michael, and Brand, Myles, eds. 1980. *Bowling Green Studies in Applied Philosophy 2: Action and Responsibility.* Bowling Green, Ohio: Bowling Green State University Press.

Brand, Myles, and Walton, Douglas, eds. 1976. *Action Theory.* Dordrecht: D. Reidel Publishing Co.

Brandt, Richard. 1958. Blameworthiness and Obligation. In *Essays in Moral Philosophy,* ed. A. I. Melden, pp. 3–39. Seattle: University of Washington Press.

Bratman, Michael E. 1987. *Intention, Plans, and Practical Reason.* Cambridge, Mass.: Harvard University Press.

Bratman, Michael E. 1996. Identification, Decision, and Treating as a Reason. *Philosophical Topics* 24:1–18.

Bratman, Michael E. 1997. Responsibility and Planning. *Journal of Ethics* 1:27–43.

Bratman, Michael E. 1999*a.* "Agency, Will, and Practical Reason: In Search of a Middle Way." Stanford, Calif.: Stanford University, Philosophy Department.

Bratman, Michael E. 1999*b. Faces of Intention: Selected Essays on Intention and Agency.* Cambridge: Cambridge University Press.

Campbell, Joseph Kiem. 1997. A Compatibilist Theory of Alternative Possibilities. *Philosophical Studies* 88:319–30.

Cauman, Leigh S.; Levi, Isaac; Parsons, Charles; and Schwartz, Robert, eds. 1983. *How Many Questions? Essays in Honor of Sidney Morgenbesser.* Indianapolis: Hackett Publishing Co.

Chisholm, Roderick. 1966. Freedom and Action. In *Freedom and Determinism,* Random House Studies in Philosophy, ed. Keith Lehrer, vol. 7, pp. 11–44. New York: Random House.

Chisholm, Roderick. 1971. Reflections on Human Agency. *Idealistic Studies* 1:36–46.

Chisholm, Roderick. 1976*a. Person and Object.* La Salle, Ill.: Open Court.

Chisholm, Roderick. 1976*b.* The Agent as Cause. In *Action Theory,* ed. Myles Brand and Douglas Walton, pp. 199–212. Dordrecht: D. Reidel Publishing Co.

Chisholm, Roderick. 1982*a.* Replies. In *Roderick M. Chisholm,* ed. R. J. Bogdan, pp. 141–65. Dordrecht: D. Reidel Publishing Co.

Chisholm, Roderick. 1982*b.* Self-Profile. In *Roderick M. Chisholm,* ed. R. J. Bogdan, pp. 1–16. Dordrecht: D. Reidel Publishing Co.

Chisholm, Roderick. 1995. Agents, Causes, and Events: The Problem of Free Will. In *Agents, Causes, and Events: Essays on Indeterminism and Free Will,* ed. Timothy O'Connor, pp. 95–100. New York: Oxford University Press.

Christman, John, ed. 1989. *The Inner Citadel: Essays on Individual Autonomy.* New York: Oxford University Press.

Christman, John. 1991. Autonomy and Personal History. *Canadian Journal of Philosophy* 21:1–24.

Clarke, Randolph. 1993. Toward a Credible Agent-Causal Account of Free Will. *Noûs* 27:191–203. Reprinted in O'Connor 1995.

Clarke, Randolph. 1994. Ability and Responsibility for Omissions. *Philosophical Studies* 73:195–208.

Clarke, Randolph. 1995. Indeterminism and Control. *American Philosophical Quarterly* 32:125–38.

Clarke, Randolph. 1996*a*. Agent-Causation and Event-Causation in the Production of Free Action. *Philosophical Topics* 24:19–48.

Clarke, Randolph. 1996*b*. Contrastive Rational Explanation of Free Choice. *Philosophical Quarterly* 46:185–201.

Clarke, Randolph. 1997. Review of *Responsibility and the Moral Sentiments,* by R. Jay Wallace. *Philosophy and Phenomenological Research* 57:230–32.

Coleman, Jules L., and Buchanan, Allen, eds. 1994. *In Harm's Way: Essays in Honor of Joel Feinberg.* Cambridge Studies in Philosophy and Law. Cambridge: Cambridge University Press.

Copp, David. 1997. Defending the Principle of Alternate Possibilities: Blameworthiness and Moral Responsibility. *Noûs* 31:441–56.

Cummins, Robert. 1979. Could Have Done Otherwise. *Personalist* 60:411–14.

Cummins, Robert. 1980. Culpability and Mental Disorder. *Canadian Journal of Philosophy* 10:207–32.

Curren, Randall R. 1989. The Contribution of *Nicomachean Ethics* III.5 to Aristotle's Theory of Responsibility. *History of Philosophy Quarterly* 6:261–77.

Davidson, Donald. 1980. *Essays on Actions and Events.* Oxford: Clarendon Press.

Deigh, John. 1995. Empathy and Universalizability. *Ethics* 105:743–63.

Della Rocca, Michael. 1998. Frankfurt, Fischer and Flickers. *Noûs* 32:99–105.

Dennett, Daniel C. 1984. *Elbow Room: The Varieties of Free Will Worth Wanting.* Cambridge, Mass.: MIT Press.

Double, Richard. 1988. Libertarianism and Rationality. *Southern Journal of Philosophy* 26:431–39.

Double, Richard. 1991. *The Non-Reality of Free Will.* New York: Oxford University Press.

Duff, R. A. 1990. *Intention, Agency, and Criminal Liability.* Cambridge, Mass.: Blackwell Publishers.

Duggan, Timothy, and Gert, Bernard. 1979. Free Will as the Ability to Will. *Noûs* 13:197–217. Reprinted in Fischer, ed. 1986.

Dworkin, Gerald. 1970. Acting Freely. *Noûs* 4:367–83.

Dworkin, Gerald. 1988. *The Theory and Practice of Autonomy.* New York: Cambridge University Press.

Ekstrom, Laura Waddell. 1998. Protecting Incompatibilist Freedom. *American Philosophical Quarterly* 35:281–91.

Eshleman, Andrew. 1998. "Responsibility for Character." Ph.D. dissertation, University of California, Riverside.

Feinberg, Joel. 1970*a*. Action and Responsibility. In *Doing and Deserving: Essays in the Theory of Responsibility,* ed. Joel Feinberg, pp. 119–51. Princeton, N.J.: Princeton University Press.

Feinberg, Joel. 1970*b*. Causing Voluntary Actions. In *Doing and Deserving: Essays in the Theory of Responsibility,* ed. Joel Feinberg, pp. 152–86. Princeton, N.J.: Princeton University Press.

Feinberg, Joel. 1970*c*. *Doing and Deserving: Essays in the Theory of Responsibility.* Princeton, N.J.: Princeton University Press.

Feinberg, Joel. 1986. *Harm to Self.* Vol. 3 of *The Moral Limits of the Criminal Law.* New York: Oxford University Press.

Fingarette, Herbert. 1972. *The Meaning of Criminal Insanity.* Berkeley: University of California Press.

Fischer, John Martin. 1982. Responsibility and Control. *Journal of Philosophy* 79:24–40. Reprinted in Fischer, ed. 1986.

Fischer, John Martin. 1983. Incompatibilism. *Philosophical Studies* 43:127–37.

Fischer, John Martin. 1985/86. Responsibility and Failure. *Proceedings of the Aristotelian Society* 86:251–70. Reprinted in French 1991.

Fischer, John Martin. 1986. Introduction: Responsibility and Freedom. In *Moral Responsibility,* ed. John Martin Fischer, pp. 9–61. Ithaca, N.Y.: Cornell University Press.

Fischer, John Martin, ed. 1986. *Moral Responsibility.* Ithaca, N.Y.: Cornell University Press.

Fischer, John Martin. 1987. Responsiveness and Moral Responsibility. In *Responsibility, Character and the Emotions: New Essays in Moral Philosophy,* ed. Ferdinand Schoeman, pp. 81–106. Cambridge: Cambridge University Press.

Fischer, John Martin. 1994. *The Metaphysics of Free Will: An Essay on Control.* Aristotelian Society Monograph Series, vol. 14. Cambridge, Mass.: Blackwell Publishers.

Fischer, John Martin. 1995. Libertarianism and Avoidability: A Reply to Widerker. *Faith and Philosophy* 12:119–25.

Fischer, John Martin. 1996. Review of *Responsibility and the Moral Sentiments,* by R. Jay Wallace. *Ethics* 106:850–53.

Fischer, John Martin. 1997. Responsibility, Control, and Omissions. *Journal of Ethics* 1:45–64.

Fischer, John Martin. In press. Responsibility and Self-Expression. *Journal of Ethics.*

Fischer, John Martin, and Ravizza, Mark. 1991. Responsibility and Inevitability. *Ethics* 101:258–78.

Fischer, John Martin, and Ravizza, Mark. 1992*a*. Responsibility, Freedom, and Reason. Critical review of *Freedom within Reason,* by Susan Wolf. *Ethics* 102:368–89.

Fischer, John Martin, and Ravizza, Mark. 1992*b*. The Inevitable. *Australasian Journal of Philosophy* 70:388–404.

Fischer, John Martin, and Ravizza, Mark. 1993. Introduction. In *Perspectives on Moral Responsibility,* ed. John Martin Fischer and Mark Ravizza, pp. 1–41. Ithaca, N.Y.: Cornell University Press.

Fischer, John Martin, and Ravizza, Mark, eds. 1993. *Perspectives on Moral Responsibility.* Ithaca, N.Y.: Cornell University Press.

Fischer, John Martin, and Ravizza, Mark. 1994*a*. Responsibility and History. In *Midwest Studies in Philosophy 19: Philosophical Naturalism,* ed. Peter A. French, Theodore E. Uehling, Jr., and Howard K. Wettstein, pp. 430–51. Notre Dame, Ind.: University of Notre Dame Press.

Fischer, John Martin, and Ravizza, Mark. 1994*b*. Responsibility for Consequences. In *In Harm's Way: Essays in Honor of Joel Feinberg,* ed. Jules L. Coleman and Allen Buchanan, pp. 183–208. Cambridge: Cambridge University Press. Reprinted in Fischer and Ravizza, eds. 1993.

Fischer, John Martin, and Ravizza, Mark. 1998. *Responsibility and Control: A Theory of Moral Responsibility.* Cambridge Studies in Philosophy and Law. Cambridge: Cambridge University Press.

Frankfurt, Harry G. 1969. Alternate Possibilities and Moral Responsibility. *Journal of Philosophy* 66:829–39. Reprinted in Fischer, ed. 1986.

Frankfurt, Harry G. 1971. Freedom of the Will and the Concept of a Person. *Journal of Philosophy* 68:5–20. Reprinted in Fischer, ed. 1986.

Frankfurt, Harry G. 1975. Three Concepts of Free Action: II. *Proceedings of the Aristotelian Society,* suppl. vol. 49, pp. 113–25. London: Methuen & Co. Reprinted in Fischer, ed. 1986.

Frankfurt, Harry G. 1983. What We Are Morally Responsible For. In *How Many Questions? Essays in Honor of Sidney Morgenbesser,* ed. Leigh S. Cauman, Isaac Levi, Charles Parsons, and Robert Schwartz, pp. 321–35. Indianapolis: Hackett Publishing Co. Reprinted in Fischer and Ravizza, eds. 1993.

Frankfurt, Harry G. 1987. Identification and Wholeheartedness. In *Responsibility, Character, and the Emotions,* ed. Ferdinand Schoeman, pp. 27–45. Cambridge: Cambridge University Press. Reprinted in Fischer and Ravizza, eds. 1993.

Frankfurt, Harry G. 1992. The Faintest Passion. *Proceedings and Addresses of the American Philosophical Association,* vol. 66, pp. 5–16. Newark, Del.: American Philosophical Association.

Frankfurt, Harry G. 1994. An Alleged Asymmetry between Actions and Omissions. *Ethics* 104:620–23.

French, Peter A., ed. 1972. *Individual and Collective Responsibility: Massacre at My Lai.* Issues in Contemporary Ethics. Cambridge, Mass.: Schenkman Publishing Co.

French, Peter A. 1984. *Collective and Corporate Responsibility.* New York: Columbia University Press.

French, Peter A., ed. 1991. *The Spectrum of Responsibility.* New York: St. Martin's Press.

French, Peter A., 1992. *Responsibility Matters.* Lawrence: University Press of Kansas.

French, Peter A.; Uehling, Theodore E., Jr.; and Wettstein, Howard K., eds. 1985. *Midwest Studies in Philosophy 10: Studies in the Philosophy of Mind.* Minneapolis: University of Minnesota Press.

French, Peter A.; Uehling, Theodore E., Jr.; and Wettstein, Howard K., eds. 1994. *Midwest Studies in Philosophy 19: Philosophical Naturalism.* Notre Dame, Ind.: University of Notre Dame Press.

Friedman, Marilyn. 1986. Autonomy and the Split-Level Self. *Southern Journal of Philosophy* 24:19–35.

Ginet, Carl. 1966. Might We Have No Choice? In *Freedom and Determinism,* Random House Studies in Philosophy, ed. Keith Lehrer, vol. 7, pp. 87–104. New York: Random House.

Ginet, Carl. 1980. The Conditional Analysis of Freedom. In *Time and Cause: Essays Presented to Richard Taylor, Philosophical Studies* Series in Philosophy, ed. Peter Van Inwagen, vol. 19, pp. 171–86. Dordrecht: D. Reidel Publishing Co.

Ginet, Carl. 1983. In Defense of Incompatibilism. *Philosophical Studies* 44:391–400.

Ginet, Carl. 1989. Reasons Explanation of Action: An Incompatibilist Account. *Philosophical Perspectives* 3:17–46. Reprinted in Ginet 1990; and O'Connor 1995.

Ginet, Carl. 1990. *On Action.* Cambridge: Cambridge University Press.

Ginet, Carl. 1996. In Defense of the Principle of Alternative Possibilities: Why I Don't Find Frankfurt's Argument Convincing. *Philosophical Perspectives* 10:403–17.

Ginet, Carl. 1997. Freedom, Responsibility, and Agency. *Journal of Ethics* 1:85–98.

Glannon, Walter. 1995. Responsibility and the Principle of Possible Action. *Journal of Philosophy* 92:261–74.

Glannon, Walter. 1997. Sensitivity and Responsibility for Consequences. *Philosophical Studies* 75:5–24.

Glover, Jonathan. 1970. *Responsibility,* International Library of Philosophy and Scientific Method. London: Routledge & Kegan Paul.

Goetz, Stewart. In press. Stumping for Widerker. *Faith and Philosophy.*

Greenspan, Patricia. 1978. Behavior Control and Freedom of Action. *Philosophical Review* 87:225–40. Reprinted in Fischer, ed. 1986.

Haji, Ishtiyaque. 1992. A Riddle Regarding Omissions. *Canadian Journal of Philosophy* 22:485–502.

Haji, Ishtiyaque. 1993. Alternative Possibilities, Moral Obligation, and Moral Responsibility. *Philosophical Papers* 22:41–50.

Haji, Ishtiyaque. 1998. *Moral Appraisability.* New York: Oxford University Press.

Heinaman, Robert. 1986. Incompatibilism without the Principle of Alternative Possibili-

ties. *Australasian Journal of Philosophy* 64:266–76. Reprinted in Fischer and Ravizza, eds. 1993.

Honderich, Ted, ed. 1973. *Essays on Freedom of Action.* London: Routledge & Kegan Paul.

Honderich, Ted. 1988*a. Mind and Brain.* Vol. 1 of *A Theory of Determinism.* Oxford: Clarendon Press.

Honderich, Ted. 1988*b. The Consequences of Determinism.* Vol. 2 of *A Theory of Determinism.* Oxford: Clarendon Press.

Horgan, Terence. 1979. 'Could,' Possible Worlds, and Moral Responsibility. *Southern Journal of Philosophy* 17:345–58.

Horgan, Terence. 1985. Compatibilism and the Consequence Argument. *Philosophical Studies* 47:403–11.

Howard-Snyder, Daniel, and Jordan, Jeff, eds. 1996. *Faith, Freedom, and Rationality: Philosophy of Religion Today.* Lanham, Md.: Rowman & Littlefield.

Hunt, David. 1996. Frankfurt Counterexamples: Some Comments on the Widerker-Fischer Debate. *Faith and Philosophy* 13:395–401.

Hunt, David. In press. Moral Responsibility and Unavoidable Action. *Philosophical Studies.*

Hunt, David. 1999. Freedom, Foreknowledge, and Frankfurt. Working paper. Whittier, Calif.: Whittier College, Philosophy Department.

Irwin, Terence. 1980. Reason and Responsibility in Aristotle. In *Essays on Aristotle's Ethics,* Major Thinker Series, ed. Amélie Oksenberg Rorty, vol. 2, pp. 117–56. Berkeley: University of California Press.

Kane, Robert. 1985. *Free Will and Values.* SUNY Series in Philosophy. Albany: SUNY Press.

Kane, Robert. 1994. Free Will: The Elusive Ideal. *Philosophical Studies* 75:25–60.

Kane, Robert. 1996. *The Significance of Free Will.* New York: Oxford University Press.

Kapitan, Tomis. 1996. Modal Principles in the Metaphysics of Free Will. *Philosophical Perspectives* 10:419–45.

Lehrer, Keith, ed. 1966. *Freedom and Determinism.* Random House Studies in Philosophy, vol. 7. New York: Random House.

Lehrer, Keith. 1976. 'Can' in Theory and Practice: A Possible Worlds Analysis. In *Action Theory,* ed. Myles Brand and Douglas Walton. Dordrecht: D. Reidel Publishing Co.

Lehrer, Keith. 1980. Preferences, Conditionals, and Freedom. In *Time and Cause: Essays Presented to Richard Taylor, Philosophical Studies* Series in Philosophy, ed. Peter Van Inwagen, vol. 19, pp. 187–201. Dordrecht: D. Reidel Publishing Co.

Lehrer, Keith. 1997*a.* Freedom, Preference, and Autonomy. *Journal of Ethics* 1:3–25.

Lehrer, Keith. 1997*b. Self-Trust.* Oxford: Oxford University Press.

Lewis, David K. 1981. Are We Free to Break the Laws? *Theoria* 47:113–21.

Locke, Don. 1975. Three Concepts of Free Action: I. *Proceedings of the Aristotelian Society,* suppl., 49:95–112. Reprinted in Fischer, ed. 1986.

Locke, John. [1690] 1965. *An Essay concerning Human Understanding,* abridged and ed. Maurice Cranston. New York: Collier.

May, Larry. 1992. *Sharing Responsibility.* Chicago: University of Chicago Press.

McIntyre, Alison. 1994. Compatibilists Could Have Done Otherwise: Responsibility and Negative Agency. *Philosophical Review* 103:453–88.

McKenna, Michael. 1997. Alternative Possibilities and the Failure of the Counterexample Strategy. *Journal of Social Philosophy* 28:71–85.

McKenna, Michael. 1998*a.* Does Strong Compatibilism Survive Frankfurt Counterexamples? *Philosophical Studies* 91:259–64.

McKenna, Michael. 1998*b.* The Limits of Evil and the Role of Moral Address. *Journal of Ethics* 2:123–42.

Melden, A. I., ed. 1958. *Essays in Moral Philosophy.* Seattle: University of Washington Press.

Mele, Alfred R. 1995. *Autonomous Agents: From Self-Control to Autonomy.* New York: Oxford University Press.

Mele, Alfred R. 1996. Soft Libertarianism and Frankfurt-Style Scenarios. *Philosophical Topics* 24:123–41.

Mele, Alfred, and Robb, David. 1998. Rescuing Frankfurt-Style Cases. *Philosophical Review* 107:97–112.

Meyer, Susan Sauvé. 1993. *Aristotle on Moral Responsibility: Character and Cause.* Issues in Ancient Philosophy. Oxford: Blackwell.

Morris, Herbert. 1976. *On Guilt and Innocence: Essays in Legal Philosophy and Moral Psychology.* Berkeley: University of California Press.

Murphy, Jeffrie. 1972. Moral Death: A Kantian Essay on Psychopathy. *Ethics* 82:284–98.

Nagel, Thomas. 1995. The Problem of Autonomy. Excerpted from chap. 7 of Thomas Nagel, *The View from Nowhere* (New York: Oxford University Press, 1986). In *Agents, Causes, and Events: Essays on Indeterminism and Free Will,* ed. Timothy O'Connor, pp. 33–42. New York: Oxford University Press.

Nathan, N. M. L. 1992. *Will and World: A Study in Metaphysics.* Oxford: Clarendon Press.

Naylor, Margery Bedford. 1984. Frankfurt on the Principle of Alternate Possibilities. *Philosophical Studies* 46:249–58.

Nozick, Robert. 1981. *Philosophical Explanations.* Cambridge, Mass.: Harvard University Press.

O'Connor, Timothy. 1993*a.* Alternative Possibilities and Responsibility. *Southern Journal of Philosophy* 31:345–72.

O'Connor, Timothy. 1993*b.* Indeterminism and Free Agency: Three Recent Views. *Philosophy and Phenomenological Research* 53:499–525.

O'Connor, Timothy, ed. 1995. *Agents, Causes, and Events: Essays on Indeterminism and Free Will.* New York: Oxford University Press.

O'Connor, Timothy. 1996. Why Agent Causation? *Philosophical Topics* 24:143–51.

O'Connor, Timothy. In press. Causation and the Will: A Metaphysical Study of the Basis of Responsibility. New York: Oxford University Press.

Oshana, Marina. 1997. Ascriptions of Responsibility. *American Philosophical Quarterly* 34:71–83.

Oshana, Marina. 1998. Personal Autonomy and Society. *Journal of Social Philosophy* 29:81–102.

Otsuka, Michael. 1998. Incompatibilism and the Avoidability of Blame. *Ethics* 108:685–701.

Pereboom, Derk. 1999. Living without Free Will. Burlington, Vt.: University of Vermont, Philosophy Department.

Ravizza, Mark. 1994. Semi-compatibilism and the Transfer of Non-Responsibility. *Philosophical Studies* 75:61–93.

Reid, Thomas. 1852. *Essays on the Active Powers of the Human Mind.* In *The Works of Thomas Reid, D.D.,* ed. Sir William Hamilton. 3d ed., vol. 1. Edinburgh: Maclachlan & Stewart.

Roberts, Jean. 1989. Aristotle on Responsibility for Action and Character. *Ancient Philosophy* 9:23–36.

Rorty, Amélie Oksenberg, ed. 1976. *The Identities of Persons.* Berkeley: University of California Press.

Rorty, Amélie Oksenberg, ed. 1980. *Essays on Aristotle's Ethics.* Major Thinkers Series, vol. 2. Berkeley: University of California Press.

Rowe, William L. 1989. Causing and Being Responsible for What Is Inevitable. *American Philosophical Quarterly* 26:153–59. Reprinted in Fischer and Ravizza, eds. 1993.

Rowe, William L. 1991. *Thomas Reid on Freedom and Morality.* Ithaca, N.Y.: Cornell University Press.

Russell, Paul. 1995. *Freedom and Moral Sentiment: Hume's Way of Naturalizing Responsibility.* New York: Oxford University Press.

Schlossberger, Eugene. 1992. *Moral Responsibility and Persons.* Philadelphia: Temple University Press.

Schoeman, Ferdinand, ed. 1987. *Responsibility, Character, and the Emotions: New Essays in Moral Psychology.* Cambridge: Cambridge University Press.

Schopp, Robert. 1991. *Automatism, Insanity, and the Psychology of Criminal Responsibility.* Cambridge Studies in Philosophy and Law. Cambridge: Cambridge University Press.

Shatz, David. 1985. Free Will and the Structure of Motivation. In *Midwest Studies in Philosophy 10: Studies in the Philosophy of Mind,* ed. Peter A. French, Theodore E. Uehling, Jr., and Howard K. Wettstein, pp. 451–82. Minneapolis: University of Minnesota Press.

Shatz, David. 1988. Compatibilism, Values, and 'Could Have Done Otherwise.' *Philosophical Topics* 16:151–200.

Sinnott-Armstrong, Walter. 1988. *Moral Dilemmas.* Oxford: Blackwell Publishers.

Slote, Michael. 1980. Understanding Free Will. *Journal of Philosophy* 77:136–51. Reprinted in Fischer, ed. 1986.

Slote, Michael. 1982. Selective Necessity and the Free-Will Problem. *Journal of Philosophy* 79:5–24.

Speak, Daniel. 1999. Fischer and Avoidability: A Reply to Widerker and Katzoff. *Faith and Philosophy* 16:239–47.

Strawson, Galen. 1986. *Freedom and Belief.* Oxford: Clarendon Press.

Strawson, P. F. 1962. Freedom and Resentment. *Proceedings of the British Academy,* vol. 48, pp. 187–211. London: Oxford University Press. Reprinted in Watson 1982; also reprinted in Fischer and Ravizza, eds. 1993.

Stump, Eleonore. 1988. Sanctification, Hardening of the Heart, and Frankfurt's Concept of Free Will. *Journal of Philosophy* 85:395–420. Reprinted in Fischer and Ravizza, eds. 1993.

Stump, Eleonore. 1990. Intellect, Will, and the Principle of Alternate Possibilities. In *Christian Theism and the Problems of Philosophy.* Library of Religious Philosophy, ed. Michael D. Beaty, vol. 5, pp. 254–85. Reprinted in Fischer and Ravizza, eds. 1993.

Stump, Eleonore. 1995. Cartesian Dualism and Materialism with Reductionism. *Faith and Philosophy* 12:505–31.

Stump, Eleonore. 1996a. Libertarian Freedom and the Principle of Alternate Possibilities. In *Faith, Freedom, and Rationality: Philosophy of Religion Today,* ed. Daniel Howard-Snyder and Jeff Jordan, pp. 73–88. Lanham, Md.: Rowman & Littlefield.

Stump, Eleonore. 1996b. Persons: Identification and Freedom. *Philosophical Topics* 24:183–214.

Stump, Eleonore. 1997. Aquinas's Account of Freedom: Intellect and Will. *Monist* 80:576–97.

Stump, Eleonore. 1999a. Alternative Possibilities and Responsibility: The Flicker of Freedom. Paper delivered at the American Philosophical Association Pacific Division Meetings, Los Angeles, March 1998.

Stump, Eleonore. 1999b. Augustine on Free Will. In *Cambridge Companion to Augustine,* ed. Eleonore Stump and Norman Kretzmann. Cambridge: Cambridge University Press.

Stump, Eleonore. In press. Dust, Determinism, and Frankfurt: A Reply to Goetz. *Faith and Philosophy.*

Stump, Eleonore, and Kretzmann, Norman, eds. 1999. *Cambridge Companion to Augustine.* Cambridge: Cambridge University Press.

Swanton, Christine. 1992. *Freedom: A Coherence Theory.* Indianapolis: Hackett Publishing Co.

Taylor, Charles. 1976. Responsibility for Self. In *The Identities of Persons,* ed. Amélie Oksenberg Rorty, pp. 281–99. Berkeley: University of California Press.

Taylor, Richard. 1966. *Action and Purpose.* Englewood Cliffs, N.J.: Prentice-Hall.

Taylor, Richard. 1982. Agent and Patient. *Erkenntnis* 18:111–26.

Taylor, Richard, 1992. *Metaphysics.* 4th ed. Englewood Cliffs, N.J.: Prentice-Hall.

Thalberg, Irving. 1978. Hierarchical Analyses of Unfree Action. *Canadian Journal of Philosophy* 8:211–26.

Thorp, John. 1980. *Free Will: A Defense of Neurophysiological Determinism.* London: Routledge & Kegan Paul.

Tomberlin, James E., ed. 1992. *Philosophical Perspectives VI: Ethics.* Atascadero, Calif.: Ridgeview Publishing Co.

Van Inwagen, Peter. 1975. The Incompatibility of Free Will and Determinism. *Philosophical Studies* 27:185–99.

Van Inwagen, Peter. 1978. Ability and Responsibility. *Philosophical Review* 87:201–24. Reprinted in Fischer, ed. 1986.

Van Inwagen, Peter. 1980. The Incompatibility of Responsibility and Determinism. In *Bowling Green Studies in Applied Philosophy 2: Action and Responsibility,* ed. Michael Brady and Myles Brand, pp. 30–37. Bowling Green, Ohio: Bowling Green State University Press. Reprinted in Fischer, ed. 1986.

Van Inwagen, Peter, ed. 1980. *Time and Cause: Essays Presented to Richard Taylor. Philosophical Studies* Series in Philosophy, vol. 19. Dordrecht: D. Reidel Publishing Co.

Van Inwagen, Peter. 1983. *An Essay on Free Will.* Oxford: Clarendon Press.

Van Inwagen, Peter. 1993. *Metaphysics.* Dimensions of Philosophy Series. Boulder, Colo.: Westview Press.

Van Inwagen, Peter. 1994. When the Will Is Not Free. *Philosophical Studies* 75:95–113.

Van Straaten, Zak, ed. 1980. *Philosophical Subjects: Essays Presented to P. F. Strawson.* Oxford: Clarendon Press.

Vihvelin, Kadri. 1998. Critical notice of *The Metaphysics of Free Will,* by John Martin Fischer. *Noûs* 32:406–20.

Vuoso, George. 1987. Background, Responsibility, and Excuse. *Yale Law Journal* 96: 1661–86.

Wallace, James D. 1978. *Virtues and Vices.* Contemporary Philosophy. Ithaca, N.Y.: Cornell University Press.

Wallace, R. Jay. 1994. *Responsibility and the Moral Sentiments.* Cambridge, Mass.: Harvard University Press.

Warfield, Ted. 1996. Determinism and Moral Responsibility Are Incompatible. *Philosophical Topics* 24:215–26.

Watson, Gary. 1975. Free Agency. *Journal of Philosophy* 72:205–20. Reprinted in Fischer, ed. 1986.

Watson, Gary, ed. 1982. *Free Will.* Oxford Readings in Philosophy. Oxford: Oxford University Press.

Watson, Gary. 1987*a.* Free Action and Free Will. *Mind* 96:145–72.

Watson, Gary. 1987*b.* Responsibility and the Limits of Evil: Variations on a Strawsonian Theme. In *Responsibility, Character, and the Emotions: New Essays in Moral Psychology,* ed. Ferdinand Schoeman, pp. 256–86. Cambridge: Cambridge University Press. Reprinted in Fischer and Ravizza, eds. 1993.

Watson, Gary. 1996. Two Faces of Responsibility. *Philosophical Topics* 24:227–48.

Widerker, David. 1991. Frankfurt on 'Ought Implies Can' and Alternative Possibilities. *Analysis* 51:222–24.

Widerker, David. 1995*a.* Libertarian Freedom and the Avoidability of Decisions. *Faith and Philosophy* 12:113–18.

Widerker, David. 1995*b.* Libertarianism and Frankfurt's Attack on the Principle of Alternative Possibilities. *Philosophical Review* 104:247–61.

Widerker, David, and Katzoff, Charlotte. 1996. Avoidability and Libertarianism: A Response to Fischer. *Faith and Philosophy* 13:415–21.

Wiggins, David. 1973. Towards a Reasonable Libertarianism. In *Essays on Freedom of Action,* ed. Ted Honderich, pp. 31–62. London: Routledge & Kegan Paul.

Wolf, Susan. 1980. Asymmetrical Freedom. *Journal of Philosophy* 77:151–66. Reprinted in Fischer, ed. 1986.

Wolf, Susan. 1981. The Importance of Free Will. *Mind* 90:386–405. Reprinted in Fischer and Ravizza, eds. 1993.

Wolf, Susan. 1987. Sanity and the Metaphysics of Responsibility. In *Responsibility, Character, and the Emotions: New Essays in Moral Psychology,* ed. Ferdinand Schoeman, pp. 46–62. Cambridge: Cambridge University Press.

Wolf, Susan, 1990. *Freedom within Reason.* New York: Oxford University Press.

Wyma, Keith. 1997. Moral Responsibility and Leeway for Action. *American Philosophical Quarterly* 34:57–70.

Yaffe, Gideon. In press. *Liberty Worth the Name: Locke on Free Agency.* Princeton, N.J.: Princeton University Press.

Yaffe, Gideon. 1999*a*. Free Will and Agency at Its Best. Working paper. Los Angeles: University of Southern California, Department of Philosophy.

Yaffe, Gideon. 1999*b*. 'Ought' Implies 'Can' and the Principle of Alternative Possibilities. Working paper. Los Angeles: University of Southern California, Department of Philosophy.

Zimmerman, David. 1994. Acts, Omissions, and 'Semi-compatibilism.' *Philosophical Studies* 73:209–23.

Zimmerman, Michael J. 1984. *An Essay on Human Action.* New York: Peter Lang.

Zimmerman, Michael J. 1988. *An Essay on Moral Responsibility.* Totowa, N.J.: Rowman & Littlefield.