

6

Moral Responsibility for Unwitting Omissions

A NEW TRACING VIEW

Dana Kay Nelkin and Samuel C. Rickless

6.1. Introduction

It is a widespread belief that moral responsibility for an action requires satisfaction of both an epistemic condition and a control condition: in order to be morally responsible for doing A, one must be aware of doing A and of A's moral significance (epistemic condition), and one must be in control of whether A occurs (control condition). If (under normal circumstances) Miranda doesn't realize that she will alert the burglar by turning on the light, then she is not morally responsible for alerting the burglar; nor is she morally responsible for alerting the burglar if a muscle spasm (or neural stimulation by an evil scientist) causes her to flip the light switch. When we focus on the realm of those actions over which we have control, it stands to reason, then, that we lack moral responsibility for those we perform unwittingly.

Unwitting omissions, however, leave some philosophers scratching their heads. For common-sense morality holds many unwitting omitters morally responsible for their omissions (and for the consequences thereof), even though they appear to lack both awareness and control. Examples are not difficult to come by, whether it be moral responsibility for unwittingly leaving a dog trapped in a car outside on a hot day (a high-stakes case—see Sher 2009) or moral responsibility for unwittingly failing to pick up milk on one's way home from work despite having promised earlier to do so (a low-stakes case—see Clarke 2014). And yet, if moral responsibility requires

awareness of one's omission and of its moral significance, as well as control, then it would appear that the unwitting protagonists of these cases are not, in fact, morally responsible for their omissions.¹

One response to this conundrum is that common-sense morality is mistaken and we should embrace skepticism about moral responsibility for unwitting omissions: general principles governing the relation between moral responsibility and epistemic states preclude moral responsibility for (many, most, or all) unwitting omissions (see Zimmerman 1997; Rosen 2004; King 2009). This, however, is a bitter pill to swallow, and numerous theorists have attempted to achieve better consilience with common sense.

For nonskeptics, the challenge is to explain how it could be possible for the unwittingness of an omission not to defeat moral responsibility for it. Nonskeptics fall into two broad categories that map onto a distinction between derivative and nonderivative (or basic) moral responsibility. Whether moral responsibility for an action or omission is derivative or basic depends on whether it derives from moral responsibility for a previous action or omission. For example, it is commonly held that although one might not be *basically* morally responsible for having crashed one's car into a pole while driving under the influence, one would be *derivatively* morally responsible for the crash inasmuch as one was morally responsible for having become inebriated in the first place. We will use "basic theorists" to refer to those who seek to vindicate and explain moral responsibility for unwitting omissions as basic, while we will use "Tracers" to refer to those who seek to vindicate and explain moral responsibility for unwitting omissions as derivative from some previous time at which one fulfilled the basic conditions for moral responsibility.

Basic theorists include Attributionists (e.g., Smith 2005 and this volume; Talbert this volume), who tie moral responsibility for conduct to the quality of will manifested or revealed by that conduct. For example, if one's having forgotten about the dog in the car reveals an insufficient degree of care for the dog (and/or care for those to whom the dog is dear), or if one's having forgotten to purchase the milk reveals an insufficient degree of care for

1. Some cases of what might be called "unwitting actions" also challenge the epistemic condition (e.g., a case in which one didn't realize that one was telling an embarrassing story about one's friend while a tape recorder was playing, but, intuitively, one *should* have realized). And similarly, some cases of action that seem not to be in our control also challenge the control condition (e.g., the classic case of the drunk driver who risks killing pedestrians). Ideally, an account of moral responsibility will provide a unified explanation for all of these cases; unwitting omissions are an ideal test case because both awareness and control appear to be missing.

those in one's family who plan to put milk in their coffee or add milk to their cereal the next morning, then one is morally responsible for the relevant omissions; otherwise, not. Attributionists therefore deny that there is any conceptual connection between moral responsibility on the one hand and epistemic or control conditions on the other. At best, the presence of knowledge and the existence of control can provide (defeasible) evidence for the quality of will that is conceptually tied to moral responsibility.

Other basic theorists (e.g., Clarke 2014) understand the existence of moral responsibility for unwitting omissions to depend on whether or not one falls below a particular standard fixed in part by one's cognitive and volitional abilities. (Call this the "Below Standard" view.) For example, if one's cognitive abilities are such that, under ordinary circumstances, one typically retains information relevant to whether one's pets might be in danger and one typically remembers to run errands on the way home after having agreed to do so on the day, then one's failure to attend to the fact that the dog is sweltering in the hot car and one's failure to remember to stop for milk fall below the cognitive standard fixed by one's general abilities. According to the Below Standard view, when that happens one is morally responsible; otherwise, not.

In a similar vein, some take moral responsibility for omissions to depend on whether the omission was caused by the normal operation of the cognitive-affective-conative system that makes one the person one is, or rather by some glitch or fault in the system. (Call this the "Underlying Self" view, and see Sher 2009.) The thought here is that the relevant dividing line lies between "psychosis, massive brain hemorrhage, and total paralysis" on the one side, and constellations of traits that support an agent's rational capacities on the other. According to the Underlying Self view, when the former features, states, or events are causally responsible for an unwitting omission, one is absolved of responsibility for that omission; but no such absolution obtains when the very same system that typically supports one's rational behavior misfires.

Attributionists, Below-Standardists, and Underlying Self theorists all agree that *basic* moral responsibility for unwitting omissions is possible. Theoretically, securing this result requires theorists of all three types to revise or weaken the standard epistemic and/or control conditions on basic moral responsibility for actions. Their mutual disagreements, such as they are, concern what they take to be the particular jointly necessary and individually sufficient conditions for basic moral responsibility.

According to Tracers (see, e.g., Fischer and Tognazzini 2009), whether one is morally responsible for an unwitting omission depends not on facts

about one's cognitive-affective-conative system at the time of the omission (say, T₂), but on facts of this sort at some relevant *prior* time (say, T₁) and on some appropriate connection between the T₁-facts and the T₂-facts. Tracers diverge on the question of what moral responsibility for an unwitting omission at T₂ traces back to. For example, Wallace (1994), Rosen (2004), and Levy (2009) trace back to a previous action or choice at a time when the agent was aware of the possible consequences that in fact come to pass at T₂.² Others, such as FitzPatrick (2008) and Shabo (2015), trace back to one or more previous exercises of agency that do not necessarily include such awareness.

The head scratching begins in earnest when one realizes that all existing nonskeptical strategies for preserving the common-sense belief in the possibility of moral responsibility for unwitting omissions, whether basic or derivative, face a number of seemingly insurmountable objections. This pushes theorists in the direction of skepticism, which is itself contrary to common sense. Is there any way out of this conundrum? We think there is.

In section 6.2, we describe some serious objections to each of three Basic Theories and in section 6.3 to standard tracing views. Refusing to accept skepticism about moral responsibility for unwitting omissions, in section 6.4 we begin by articulating what we take to be the right *control* condition on moral responsibility, one that is weaker than the standard volitional or decisional control condition. The relevant sort of control, we argue, is fixed by having the right sort of opportunity to act in certain ways. In section 6.5, we show that a tracing view that builds in the right kind of control condition (call it the "Opportunity Tracing" view) can handle a representative sample of cases that have been offered in defense of the possibility of moral responsibility for unwitting omissions. On our view, the epistemic and control conditions are not in fact two separate conditions; rather, the control required for moral responsibility itself requires a kind of awareness. In section 6.6, we reply to various objections that might be offered against the Opportunity Tracing view. We then close by summarizing the results of our investigation.

Before beginning to canvass extant views, it is important to clarify the notion of moral responsibility we are interested in. In particular, we are interested in responsibility and blameworthiness in the "accountability" sense (see Watson 2004). When people are responsible and blameworthy in this sense, they have violated obligations and appropriate demands, and are

2. Fischer and Tognazzini (2009) take it that responsibility for an omission at T₂ traces back to an "exercise of agency" but not necessarily to a choice. We will take up this view in section 6.3.

reasonably subject to the reactive attitudes such as resentment, indignation, and guilt.³

6.2. Challenges for Basic Theories

Attributionism is the (broadly Kantian) view that moral responsibility for one's conduct (action or omission) depends entirely and solely on the quality of will manifested in that conduct. There is much to be said for such a view. In general, for example, we don't hold people morally responsible for accidents, unless they stem from lack of due care. If Gabriel is driving carefully and well within the speed limit and a stray dog suddenly runs in front of his car, he is not morally responsible for the car's hitting the dog. By contrast, if Gabriel's car hits a dog because he is driving carelessly, paying insufficient attention to his surroundings and thereby manifesting indifference to the lives of others, then he is morally responsible for the collision. Attributionism focuses on the fact that attitudes matter. Moral responsibility, at least in many cases, appears to depend on whether one's attitudes are properly oriented toward the right and the good.

But Attributionism struggles to make sense of common-sense intuitions. One such intuition is that it is possible for agents to be morally responsible (and blameworthy) for harmful conduct even when their attitudes are properly oriented. Perhaps, as some Attributionists argue (see Talbert in this volume), it is a mistake to blame good-willed people for their unwitting omissions in low-stakes cases, such as forgetting to purchase milk on the way home: in such circumstances, it may be more appropriate for the milkless to feel frustration or regret. And it may also be the case that blaming others in high-stakes cases is a symptom of attribution error, ascribing bad intentions or insufficiently caring attitudes to those who actually lack them (see Talbert in this volume). But in both low-stakes and high-stakes cases, especially in the latter, agents often and quite typically blame *themselves* for their unwitting omissions, as well as for ensuing harms. This phenomenon can't be explained as the product of attribution error, because such error involves the disposition to mistakenly ascribe certain attitudes to *others*. Moreover, as far as we are aware, there is no evidence for the claim that subjects are prone to over-*self*-attribute bad intentions or insufficient care. Indeed, anecdotal evidence suggests the opposite: people are generally disposed to judge themselves far less harshly than they are disposed to judge others.

3. On our view, these attitudes presuppose attributions of desert of blame. (See Nelkin 2013.)

Consider Sher's Alessandra, who leaves her dog locked in the car on a sweltering day (as she thinks, just for a few minutes) to pick her children up from school. When she enters the school, she is given news that distracts her ("a tangled tale of misbehavior, ill-considered punishment, and administrative bungling"), as a result of which she forgets all about her dog while she is busy sorting things out at school. As she finds herself walking out into the parking lot several hours later, it occurs to Alessandra that she has left the dog in the car. As Sher tells the story, as she approaches the car, Alessandra sees that the dog is "unconscious from heat prostration." In a more extreme version of the story, the dog is dead. We, thinking, perhaps, "There but for the grace of God go I," might feel a great deal of compassion for Alessandra. But many of us will blame her, at least inwardly, and most Alessandras will blame themselves. When the individual left in the car is a child (as sometimes happens in real life),⁴ many Alessandras find it difficult to go on living. This is not because they have suffered a devastating loss or because they suffer from deep regret. (Most parents who lose children to cancer do not feel the kind of suicidal-thought-inducing guilt that most Alessandras feel.) It's because they hold themselves to account for the omission and the harm. A few Alessandras may worry that they didn't care enough for their dog, and that it was this insufficiency of care that was responsible for their omission (to remember, to go back to the car). But few will worry that they didn't love their *child* enough, even as they blame themselves. Admittedly, many of us are not disposed to blame Alessandra (by the way, a phenomenon that is not easy to square with the attribution-error hypothesis), even when her omission results in the death of her child. But this is not because we do not find her *blameworthy*. It is, more likely, because we think, in addition to the "There but for the grace of God" thought, that blaming Alessandra would be pointless: there is no way to bring her child back, the loss cannot be processed as a learning experience, and Alessandra is already suffering more than we could possibly imagine.⁵

4. See, for example, "Infant Son Left in Car by UCI Professor Dies" (*Los Angeles Times*, August 9, 2003; <http://articles.latimes.com/2003/aug/09/local/me-deadchild9>) and "Police: Baby Dies after Being Left in Hot Car by Father" (*Des Moines Register*, June 13, 2016; <http://www.desmoinesregister.com/story/news/crime-and-courts/2016/06/12/infant-dies-unattended-car-police-arrest-father/85790806/>).

5. Another case that opponents of Attributionism find problematic for the view but that Attributionists embrace is the case in which agents' omissions reflect bad attitudes, where they are attitudes that the agents had no opportunity to avoid having or bypassing in their actions. For example, a teenager raised in a very religious community that treats homosexuality as sinful might then act from malicious attitudes in bullying a classmate. But it seems that at the least responsibility is mitigated in that case, compared to a teenager who was not

Unlike Attributionists, Below-Standardists hold that moral responsibility for an unwitting omission depends solely on whether the omission results from a substandard episode of forgetting or inattentiveness. As a proposal designed to make room for basic moral responsibility for unwitting omissions, Below-Standardism gets many cases right. And there is something intuitive in the thought of tying moral responsibility to the failure to abide by a cognitive-volitional norm. But Below-Standardism faces a serious objection.

Below-Standardists emphasize that the standard that determines whether an agent is morally responsible for an unwitting omission "is neither a moral obligation nor an ideal." Instead, the standard "is one whose application is sensitive to an individual's own cognitive and volitional capacities and to the situation that individual is in" (Clarke 2014, 167). But once one has removed any moral tint from the relevant standard, it becomes difficult to understand why moral responsibility should depend on it: why should my degree of moral responsibility for an omission be fixed by how often I forget things in this sort of situation? Consider, for example, a particularly forgetful Alessandra. From a very young age, Alessandra has been scatterbrained and disorganized, especially when distracted. When she commits thoughts to memory, she only rarely successfully recalls the information when something else in her environment becomes particularly salient. When she leaves her dog (or child) in the car, she has no good reason to believe that she will become distracted when she enters the school to pick up her children. What the Below-Standardist must say about this case is that, in forgetting all about her dog (or her child) trapped in the hot car, Alessandra does not fall below any cognitive or volitional standard fixed by her situation and her cognitive-volitional capacities, and hence is not morally responsible for her omission (to remember, to go back for the dog/child). But common sense judges otherwise: indeed, if anything, forgetful Alessandra strikes us as even more blameworthy than a significantly less forgetful counterpart.

Now consider Randy, who has promised his wife to stop by the store for milk after work on his way home, and imagine that Randy has a *prodigious* memory. In situations of this sort, he rarely forgets what he has promised

raised in such a community. The Attributionist has no way of distinguishing the two, and yet, intuitively, to our minds, at least, there is a moral difference. This is an implication many Attributionists welcome, so we do not expect this sort of case to sway them. On the other hand, it is not always clear that Attributionists are aiming to explain moral responsibility in the accountability sense. (See, for example, Björnsson [forthcoming].)

to do, no matter how minor, even when massively distracted. It is very difficult for him; but he almost always rises to the challenge. This time, though, Randy is the only witness to a near-fatal car crash on his way to the store. A child is trapped in a car that has rolled over, and Randy, Good Samaritan that he is, rushes to save her. The police arrive, and Randy patiently answers their questions. Shortly thereafter, a local TV crew shows up, and pesters him for an interview, an invitation he repeatedly declines. After extricating himself with difficulty from the scene, Randy drives straight home without remembering to stop for milk. The Below-Standardist must say that prodigious Randy is morally responsible for his omission to purchase the milk, because his failure to remember to stop at the store falls below a standard that applies to him in such situations. But common sense judges otherwise, absolving him of moral responsibility for his unwitting omission. It seems that failing to strike a very high standard that one has expended great effort to set does not thereby make one on the hook when one would have avoided culpability had one simply worked less hard in the past.

Underlying Self theorists claim that moral responsibility for an unwitting omission depends on whether the omission is caused by a fault in the constellation of cognitive-affective-conative traits that supports one's rational capacities. A major problem with this view stems from the fact that which traits happen to support which capacities is a largely contingent matter. For example, nothing rules out the possibility of one's rational capacities being supported by an extreme phobia. Suppose, then, that Alessandra has a severe and incurable form of ophidiophobia (fear of snakes) that supports her rational capacities (this being one reason why her phobia is incurable). She has worked very hard to avoid snakes, and lives in a large metropolis where snakes are very uncommon. Unbeknownst to Alessandra, however, a boy at her children's school has brought his snake in for show-and-tell. The snake escapes and slithers under the classroom door and down the hallway, just as Alessandra walks in to pick up her children. Overcome by fear, Alessandra is transfixed, unable to move. Everything that she has committed to intermediate-term memory disappears, and she forgets all about the dog (or child) in the car. According to the Underlying Self theorist, ophidiophobic Alessandra is morally responsible for her failure to remember, and for her failure to go back for, her dog (child). This is because her omission resulted from her phobia, which is a glitch in the cognitive-affective-conative system that happens to support her rational capacities. And yet this is the sort of case in which common sense absolves Alessandra of moral responsibility.

6.3. Challenges for Two Tracing Views

We have now canvassed problems faced by three different accounts of *basic* moral responsibility for unwitting omissions. These problems strongly suggest that the key to preserving common-sense intuitions in this area is to embrace the view that moral responsibility for such omissions traces back to an earlier moment in time when one met the basic conditions for responsibility, when one had the specific opportunity to do something that would have significantly reduced the chances of the later instance of unfortunate unwittingness. This is what Tracers believe, and we agree.

However, most tracing views take us back to a time when the relevant agent not only *recognizes* the very real risk of later memory or attention failure, but also *embraces* the risk by deciding to do something that will increase, or by deciding not to do something that will decrease, the probability of such failure. There is certainly a great deal to be said in favor of these “Decision Tracing” views. Suppose Peter drives his car into a storefront, and does so because his processing of visual information and ability to steer have been compromised by severe inebriation. If we learn that Peter became inebriated while driving because a vicious enemy of his had earlier spiked his drink without his knowledge, then we are not inclined to hold him morally responsible for the property damage. But if we learn that Peter earlier *chose* to imbibe a few too many alcoholic beverages while he was out partying with friends, knowing full well that too much alcohol in one’s system reduces visual acuity and depresses reaction times and that inebriated people underestimate the risks involved in driving under the influence, then we are inclined to throw the book at him. Tracing views capture these intuitions, and can make sense of many situations that would otherwise be puzzling.

Consider, for example, Sher’s Wren, “on guard duty in a combat zone” where “there is real danger.” Bored and tired, “lulled by the sound of the wind in the leaves,” Wren “has twice caught herself dozing and shaken herself awake.” Eventually, however, Wren succumbs to fatigue and “falls unto a deep slumber, leaving the compound unguarded.” Undoubtedly, Wren is morally responsible for her failure to guard the compound, whether or not the enemy takes advantage of her failure to attack it. Sher argues that Wren is a poster child for the Underlying Self view. But, as it happens, she is just as much a poster child for Decision Tracing views. For when Wren earlier caught herself dozing, she was well aware of the fact that she was in danger of nodding off, and yet did nothing about it, presumably because she decided to ignore the very real risk of falling asleep. The fact that Wren is morally responsible for her omission to guard the compound is therefore

traceable to an earlier witting decision not to take steps that would have reduced the chances that she would fall asleep, steps such as turning on the radio, lowering the ambient temperature, or arranging for a companion or replacement.

Still, as many Anti-Tracers have pointed out, Decision Tracers struggle to capture common-sense intuitions in many cases. For with respect to many unwitting omissions it is difficult to find an earlier time at which the agent actually *decided* to ignore relevant information and forgo the opportunity to prevent the later omission (see Vargas 2005; McKenna 2008; Smith 2008; King 2009; Khoury 2012; Shabo 2015). When Randy’s thoughts wander into philosophy as he is driving home from work, the thoughts he experiences are not the result of any *decision* to allow them to wander. When Alessandra walks into school to pick up her children, she does not decide to assume the risk that she will become distracted and lose track of the fact that her dog (child) is trapped in a sweltering car. And yet common sense counts Randy and Alessandra morally responsible for their unwitting omissions, an intuition that Decision Tracers find themselves unable to explain.⁶

The moral of this story is that proponents of tracing views should not hold that moral responsibility traces back to a *conscious decision*. But to what, then, does moral responsibility trace back?

A second sort of view (call it the “Akrasia-or-Vice Tracing view”) answers the question by tracing back to akrasia (weakness of will) or to previous exercises of vices. That is, the Akrasia-or-Vice Tracing view retains the idea that responsibility traces back to decisions, but rejects the idea that they must be made with a kind of foresight of the future culpable action or omission. According to FitzPatrick, “Ignorance, whether circumstantial or normative, is culpable if the agent could reasonably have been expected to take measures

6. Another consideration that has been raised as an objection to tracing views generally is that tracing is not needed to account for all we want to account for. On this view, we can limit responsibility to (1) responsibility in the basic sense for actions and omissions and to (2) the consequences of what we are basically responsible for, which might include further actions or omissions qua consequences. (See King 2009 and Agule 2016.) Addressing these arguments in any detail is beyond the scope of this chapter, but for our purposes here, we are using “tracing” in an ecumenical way so as to include the kind of view defended by King and Agule. Whether advocates of tracing views see unwitting omissions as something for which one is derivatively responsible qua consequence of the earlier action or omission for which one is basically responsible, or as something for which one is derivatively responsible in some other sense, or as neutral as between these, the only commitment we are here ascribing to tracing views is that they “trace back” responsibility for the unwitting omission to an earlier exercise of basic responsibility. So there is *some* sense in which the responsibility for unwitting omission is derivative, on tracing views as we understand them.

that would have corrected or avoided it, given his or her capabilities and the opportunities provided by the social context, but failed to do so either due to *akrasia* or due to the culpable, nonakratic exercise of such vices as overconfidence, arrogance, dismissiveness, laziness, dogmatism, incuriosity, self-indulgence, contempt, and so on" (2008, 609). Thus, if we could trace Alessandra's current lack of awareness of her dog to some earlier moment at which she exercised some vice, then she would be blameworthy on this view. Interestingly, this view combines a tracing element with an Attributionist-friendly condition. On FitzPatrick's view, it isn't that Alessandra's current omission must reflect some bad quality of will; but it must trace back to some failure of this kind.

This suggestion provides an answer to our question, but it isn't clear that it succeeds. First, though FitzPatrick mentions the opportunities provided by the social context, there is no claim that the agent must herself have an opportunity in the sense that it was accessible to *her*. Second, suppose that Alessandra exercised a vice of laziness earlier that week with no awareness that it could lead to anything as serious as risking her dog's life days from then. It isn't clear that this is sufficient for Alessandra's culpability for her unwitting omission when she leaves her dog in the car to enter the school. Is her exercise of a vice of laziness substantial enough to bear the weight of Alessandra's responsibility for the harm, or even great risk, to her dog? Perhaps the idea that she could reasonably have been expected to be aware of the risk on that day is what is really doing the work here. But then we need to know what makes it true that she could reasonably have been expected to be aware. On our view, for that to be true, she needed to have a real opportunity to be aware on that fateful afternoon, and the question remains an open one whether in order to have such an opportunity, she needed to have been aware of relevant risks at some earlier time.⁷ But this would turn the view into a tracing view that requires awareness after all.

Finally, we are intrigued by a kind of tracing view advocated by Fischer and Tognazzini (2009). According to their view, the key is that one can be responsible for an action or unwitting omission when one does not meet

7. Seth Shabo (2015) offers a tracing view that, like FitzPatrick's, does not require that the endpoint of the tracing be a conscious decision to ignore risks. On his view, a series of "morally significant" voluntary actions that cumulatively results in the later action or omission can serve as the endpoint for tracing. Here, we take it that the plausibility of the view depends on the particular kind of moral significance of the voluntary actions in question. But if at no point does the agent see, even in broad strokes, what sorts of morally bad consequences are significantly risked by the agent, then we take it that he is not responsible.

the basic conditions for responsibility at the time if it can be traced back to a past exercise of control (which might include *either* a choice to act or an omission (2009, 551). Given their earlier defense of an epistemic condition, presumably their view is that to be responsible for X at T₂ when one does not meet the basic conditions for responsibility, X must be traceable to a time T₁ when one both exercised control and was aware—at least under a broad description—of the risks of X-like acts or omissions in the future. While sympathetic to the spirit of this proposal, we have some concerns. First, we are not entirely sure what is meant by an exercise of agency that is not a choice, but rather a failure to choose. If it amounts to a choice *not* to choose X or not-X, we can understand it, but then it is at one and the same time a choice and a failure to choose (because the scope of the choice in each case is different). On the other hand, if it does not amount to a choice not to choose, we are not sure how there is an exercise of agency. On either horn of the dilemma, the view faces a challenge. It seems that the doubts raised about Decision Tracing views would transpose to the former interpretation, given that, intuitively, not all instances of culpable unwitting omissions are preceded by choice of either a first-order sort or a second-order sort; if we embrace the latter interpretation, we are left with a question of what exactly makes the failure to choose an exercise of agency in the relevant cases. For some failures to choose are clearly not exercises of agency. Benny failed to choose to water the flowers while asleep last night, but that failure was not an exercise of agency, for example. So we need to know what makes the end nodes of tracing that are failures to choose exercises of agency in cases in which agents are responsible for later omissions.

Our suggestion is that a prior exercise of agency is not needed at all; all that is needed is that the agent *had* control. But to make this plausible, we believe, contra Fischer and Tognazzini, who offer an ecumenical and nonspecific account of control, that we need a substantive account of control in order to assess the claim. In the next section, we attempt to provide such an account, and in section 6.5 to apply it to the case of unwitting omissions.⁸

8. We also diverge from Fischer and Tognazzini in that they suggest that the epistemic condition can be fulfilled by its being true that one should have been aware. On our view, that one should have been aware is itself made true only if one had the opportunity to be aware, and, in turn, this requires some earlier time at which one was in fact aware that one was risking future ignorance. Relatedly, as we explain, we take it that there is no separate epistemic condition; control itself requires awareness.

6.4. Basic Responsibility and Control

In general, basic moral responsibility for an action or omission requires control. But how should control be analyzed? There are numerous competing accounts of the kind of control needed for moral responsibility. Some (e.g., O'Connor 2016) understand control in terms of the ability to choose among alternative courses of action. Others (e.g., Clarke and Capes 2015) view it as a matter of being the source of one's actions/omissions. Yet others (e.g., Fischer and Ravizza 1998) analyze control as a matter of one's actions being caused by a moderately reasons-responsive mechanism of one's own. We do not have the space to consider and criticize these and other alternatives here. Instead, we will articulate what we take to be the best independently motivated account of the sort of action-control required for moral responsibility, to see whether it can show us the way toward a better account of moral responsibility for unwitting omissions.

Control of actions, we suggest, requires the interference-free ability to do the right thing for the right reasons (see Wolf 1990; Nelkin 2011). The ability to A is here understood to involve possession of the capacities, skills, talents, and knowledge required to A. Someone has the ability to ride a bicycle when she possesses sufficient skills and know-how to maintain her balance on two wheels while pedaling. An agent's ability to A is interference-free when nothing actually prevents her from exercising it, that is, when she has the opportunity to exercise it, which means that the situation is sufficiently amenable. Put slightly differently, then, an agent is in control of an action in the sense required for responsibility when she has the opportunity to do the right thing for the right reasons. In turn, having the opportunity means that one must be able to employ one's skills, talents, and so on in the actual situation. Thus, one must have the relevant skills and competence on the one hand, and also the cooperation of the situation on the other.⁹

This account of control is asymmetrical with respect to the ability to do otherwise. When one does the right thing for the right reasons, control does not require the ability to do otherwise (that is, the ability to do the wrong thing, or the ability to do the right thing for the wrong reasons). But when one does the wrong thing (or does the right thing for the wrong reasons), control does require the ability to do otherwise.

9. The reason for this is that as we understand "opportunity," it entails "ability."

Applied to omissions, this account of control tells us that an agent has control over a wrongful omission to A when she has the opportunity to A for the right reasons. But if control is required for basic moral responsibility, then it becomes impossible for any agent to be morally responsible in a *basic way* for her unwitting omissions. The reason for this is that the unwitting agent's lack of awareness of what she is failing to do deprives her of the *opportunity* (even if not of the *ability*) to do otherwise. Caught in the middle of a tangled tale of misbehavior, punishment, and bungling, Alessandra has completely forgotten that she earlier left her dog in a hot car without means of escape. And the fact that it does not so much as occur to her that her dog might be in trouble means that she has lost the opportunity to rescue her dog. If moral responsibility requires control, and control (in this case) requires the opportunity to do otherwise, then Alessandra is not *basically* morally responsible for her failure to go back to the parking lot and liberate her dog. For good measure, Alessandra is also not morally responsible for her failure to remember that her dog is trapped in a car under a hot sun, for the fact that she is seriously occupied and distracted deprives her of the opportunity to remember.

It stands to reason, then, that agents cannot be nonderivatively morally responsible for their unwitting omissions. There is therefore no way to account for moral responsibility in unwitting omission cases other than by tracing every such omission back to an earlier time when the requirements for moral responsibility were fulfilled.

We should note that mere control, or the bare opportunity to do the right thing for the right reasons, while necessary and sufficient on our view for basic responsibility, is not always sufficient for blameworthiness and praiseworthiness. Rather, in our view, to be blameworthy, one must have a sufficiently high quality—or, to borrow familiar language from legal theorists—a *fair* opportunity to do the right thing for the right reasons.¹⁰ To see this, consider classic cases of duress. They suggest that we might be excused for wrongdoing even though we *could* have refrained from it; as long as the situation made it simply too hard to reasonably expect us to comply with the right reasons, we would be excused. For example, if someone threatens to harm one's child unless one commits a serious crime, say, it is natural to think that one is excused, despite having the opportunity to do the right thing. At the least, one's blameworthiness can be mitigated when the quality of one's opportunity is low, as in this

10. See Brink and Nelkin (2012) for development of this idea.

case. This suggests that there is a threshold level of *quality* of opportunity for one's being blameworthy. In addition, it also gives us a way of understanding degrees of blameworthiness. The easier it is to do the right thing for the right reasons, the more blameworthy one is for failing.¹¹

6.5. *The Opportunity Tracing View*

According to the view we wish to defend, whether an agent is morally responsible for an unwitting omission at time T₂ depends entirely and solely on whether there was a prior time, T₁, at which the agent had the opportunity to do something that, as she reasonably believed, would significantly raise the likelihood of avoiding later omission. Call this the "Opportunity Tracing" view. The rationale for the view is two-pronged. The first prong is the principle that derivative moral responsibility must trace back to something for which the agent is basically morally responsible. This principle, together with the account of control necessary for moral responsibility outlined in the previous section, explains why, if an agent is to be responsible for an omission at T₂, she must have had an opportunity to act preventatively at T₁. The second prong is that awareness of a kind is necessary for having such an opportunity: in order for an agent's moral responsibility for an omission at T₂ to trace back to T₁, the agent must have good reason to believe at T₁ that there is something she can do then that will sufficiently raise the likelihood of avoiding omission at T₂.¹²

Let us see how well the Opportunity Tracing view can account for cases that have been marshaled in support of the idea that there can be basic moral responsibility for unwitting omissions. Begin with the cases of Randy and Alessandra.

Just before leaving work at 5:00 p.m., Randy promises his wife that he will stop for milk on his way home. While driving, his mind wanders and he

begins to think about the book on which he has been working. Philosophical thoughts and arguments run through his mind, theories are developed, objections parried, counterexamples imagined, and hypothetical cases tweaked. Distracted while driving, Randy forgets all about stopping for milk and arrives home empty-handed. Is he morally responsible for having omitted to stop for milk on his way home?

Yes. To understand why, one needs to focus, not on the time when Randy is thinking about his book while driving home (say, around 5:20 p.m.), not on the time when he made the promise to his wife (around 5:00 p.m.), but on the moment in between (say, at 5:10 p.m.), just before his mind begins to wander. At that time, Randy knows that he has promised that he will get milk later, so, as he knows, he has incurred an obligation. He has committed the matter to episodic memory, and this is the mechanism of recall on which he is relying as he is driving. However, as Randy also knows, episodic memory is fallible and most likely to break down when one's mind is distracted by something that takes it off one's present task. Randy has, no doubt, turned on his car radio in the past and found himself, in relatively short order, completely immersed in a news story on *All Things Considered* or laughing uproariously at jokes on *Wait, Wait, Don't Tell Me*. He knows that the radio can be so captivating that it is dangerous to turn it on if one is driving to a particularly important event (e.g., a wedding) at which late arrival would be widely viewed as a serious faux pas or insult. So Randy, wisely, doesn't turn on the radio. However, he *does* start to focus on something other than his present task, and, as he knows, this poses just as much of a risk of distraction that turning on the radio would pose. He has, no doubt, forgotten all about other tasks in the past while engaged in deep philosophical thought: "Oops, I forgot that I was supposed to wash the dishes" or "Oops, I forgot that I was supposed to deposit a check at the credit union." Aware of the fact that his mind is beginning to stray from its present focus (on driving to the store to purchase the milk), and aware of the fact that the way to the store is similar (up to a point) to the way home (which is burned by habit into his brain as one of several "rat-runs"), Randy knows that by straying into the kind of intellectual activity that demands intense focus he is running a not-insignificant risk that he will forget about his task and simply fall into the "home" rat-run by inveterate habit. At this point, he has the opportunity to raise the chances of remembering by taking some simple and easy steps: he could tell himself that he should keep his mind on task instead of getting sucked into deep philosophical thoughts (and return to his focus on driving to the store), or he could set his phone alarm to beep just around the time he expects to stop at the store (thereby shocking him out of

11. The account also points the way to an account of degrees of praiseworthiness, though that is not our focus here. For more on praiseworthiness, see Nelkin (2016). Importantly, we do not take it that control is the sole factor in quality of opportunity when assessing either blameworthiness or praiseworthiness, because quality of opportunity also tracks additional factors such as amount of sacrifice required to do the right thing.

12. We believe that this view shares a key feature with that of Carl Ginet (2000), who, following Holly Smith's (1983) introduction of the term, "benighting act," explains that culpable action (or omission) must depend on an earlier benighting act in which one was aware of relevant risks (or on an earlier act which itself depended in the right way on an earlier benighting act). We diverge from Ginet in that he takes it that responsible action also requires that one could have done otherwise either at the time or at the time of the benighting act.

his philosophical reverie). The problem is that Randy does neither of these things. Instead, he allows his mind to wander while doing nothing to prevent it from getting completely sucked into the kind of mental activity that will likely distract him from his obligation. It is for this reason that Randy is responsible for his omission to stop for milk on his way home from work.

Clarke (2014, 165; see also this volume), putting himself in Randy's shoes, counters that using a sticky note or phone alarm immediately after having made the promise to his wife "would have bordered on compulsion," and that it would have been "obsessive" for him to keep asking himself while driving whether he was forgetting something. In addition, Clarke notes that he "commonly" thinks about work "without ill effect" when he is driving. All of these remarks are true, but they do not address the main reason for thinking that he was responsible for failing to stop for milk. Given that he made the promise to his wife as he was walking out of his office, it was reasonable for him to entrust the promised task to memory *at that time*. Before it even occurred to him that he might start to think about his book or become distracted in some other way while driving, it would indeed have been bizarrely compulsive for him to have entered a reminder in his phone. As long as his mind wasn't in danger of getting distracted, it would indeed have been obsessive for him to keep reminding himself every few seconds while driving about picking up milk at the store. After all, if he is focused on his task, there is no need for constant reminders. The problem arises when, as he well knows, he is about to engage in activity that poses a serious risk of distraction. At *that* time, he has the opportunity to keep his mind on task or take anticipatory precautions against getting distracted. If he has this opportunity and, as a result, omits to discharge a moral obligation, then he is morally responsible, and potentially blameworthy, for his omission.

It must be granted, of course, that Clarke often thinks about work without ill effect when he is driving. But particular circumstances here matter greatly. When Clarke doesn't have any moral obligation to discharge while driving (apart from keeping alert and following the rules of the road), it is obvious that the likelihood that thinking about work will have some ill effect is minuscule. But *in circumstances in which it is important for him while driving not to forget to accomplish some task*, thinking about work *can* have an ill effect. Examples of this sort are legion. Suppose, for example, that Jane has thirty exams left to grade before a midnight deadline, but that she has been working diligently to crack some philosophical nut for days and the beginning of a solution suddenly occurs to her. She knows that if she starts working to crack the nut, there is a significant chance that she will get sucked in and forget about the grading she needs to do. She has the opportunity to

avoid getting sucked in (or the opportunity to set the kind of reminder that would get her back to grading in time to complete it by the deadline). So if she gets sucked in without setting the reminder, she will be morally responsible for her later failure to send her grades in before midnight.

But given how often in the past Clarke has successfully regained his focus while driving even after having allowed himself to become mentally distracted, isn't it reasonable for him to rely on the fact that his memory will kick in as he nears the turn that will take him to the store? Well, it's not particularly unreasonable, but relying on one's memory *in such circumstances* involves taking a not-insignificant risk that one will fail to discharge an obligation. Perhaps Clarke thinks that thinking about his book is important enough that he is prepared to take the risk. In that case, the relative importance of thinking about the book might go some way toward excusing, or might diminish the degree of *blameworthiness* for, his later lapse. But it does not erase his moral responsibility for the lapse.

What of Alessandra? Alessandra's situation differs somewhat from Randy's. As she leaves her car to enter the school to pick up her children, Alessandra, unlike Randy, doesn't have a particularly strong reason to think that she may become embroiled in something that will distract her sufficiently that she will fail to recall what it is important for her to remember. Her children, we may suppose, are well-behaved, their teachers reasonable and even-tempered, the school administrators competent and well-intentioned. However, the day is hot, and, as Alessandra well knows, the inside of a car can reach very high temperatures if it is left to bake in the sun for any longer than a few minutes, dogs in locked cars with the windows closed do not have the ability to escape unless freed by some external force, and (we may suppose) it is unlikely that a passerby would, if the dog were trapped, notice that the dog was in distress and take sufficient steps to free it. So, on the one hand, as she locks and leaves her car, it is reasonable for Alessandra to believe that she won't be gone longer than a few minutes; on the other hand, there is little room for error. And when the consequences of error would be devastating, there arises a stringent duty to take whatever steps would be required to avoid error.

Imagine, for example, that Alvaro is an important member of the president's Secret Service detail. His job is to keep the president safe at all times. Right now, the president is working a rope line at a rally. There are hundreds of people along the rope line hoping for the chance to shake her hand. Alvaro is next to the president, checking for signs of suspicious behavior. His iWatch vibrates. Stealing a quick glance at it, he sees that he has a text message from his son, who is at a restaurant with some of his

high school friends. Alvaro knows that he could take a few seconds to read the message on the watch (his son's texts can be insistent), and that if he does so his attention will be drawn away from the rope line. Suppose, now, that Alvaro stops for a second to read the text message, and, just at that moment, a person in the rope line lunges at the president. What are the chances that a lunge would happen at this very moment? Minuscule. And yet Alvaro's duty is to maintain total and uninterrupted focus on the people at the rope line.

Similarly, once Alessandra leaves her car knowing that the dog is trapped therein under the baking sun, she acquires a stringent duty to maintain total and uninterrupted focus on her dog's welfare, even while she is doing other things. She needs to be prepared for any eventuality, including any unanticipated interruption of her routine. If she doesn't keep her dog at the forefront of her mind as she enters the school, prepared to go back to the car at any sign that her regular routine will be interrupted, she will be morally responsible for any later failure to return to the car (unless, of course, something prevents her from returning or she acquires an even more stringent duty to stay where she is). At the moment she locks and leaves the car, Alessandra has the opportunity to maintain mental focus on her dog, the kind of mental focus that would trigger recognition of the importance of going back to the car in case of potential distraction. It is her failure to maintain that mental focus that grounds her moral responsibility for her later failure to return to the car when she learns the tangled tale of misbehavior, punishment, and bungling.

In this case, unlike in Randy's case, it seems inappropriate to suggest that constant mental reminders ("Don't forget about the dog—remember the dog—dog—dog—dog," and so on) would be out of place or obsessive. When one has a stringent duty to maintain a certain level of attention and focus, one should do whatever is required to keep that focus, even if that involves what would otherwise be considered obsessive behavior.

We conclude that the examples of Randy (a low-stakes case) and Alessandra (a high-stakes case) support the Opportunity Tracing view. What accounts for the fact that these agents are morally responsible for their unwitting omissions is that there was a time, prior to the omissions themselves, when they had the opportunity to take measures that (as they know or justifiably believe) would significantly raise the probability of avoiding such omissions, and yet they did not avail themselves of the relevant opportunity. In these sorts of cases, moral responsibility traces back, not (or not necessarily) to a *decision* to disregard information or to undertake a risky course of conduct, but to a previous *witting* failure of some sort.

6.6. *Objections and Replies*

One objection that might be raised against the Opportunity Tracing view is that it ascribes moral responsibility to agents who, intuitively, are not morally responsible for their unwitting omissions. Aren't the view's individually necessary and jointly sufficient conditions for moral responsibility so easy to fulfill that almost everyone who might be thought morally responsible for an unwitting omission turns out to be so? Consider the case in which normal Randy (not prodigious Randy) comes upon the scene of a car crash on his way home from work, and saves a child who is trapped in a car that has rolled over. Distracted, with his attention focused on the welfare of the child and with his body surging with adrenaline, Randy drives straight home having completely forgotten about his promise to stop for milk. Before his unwitting omission, wasn't there a time at which he had the opportunity to take precautions against being distracted by a later car accident? And doesn't the Opportunity Tracing view therefore entail that Randy is morally responsible for his failure to stop for milk, even in this case? And isn't that itself counterintuitive?

In reply, we certainly agree that it would be counterintuitive to hold Randy morally responsible for his unwitting omission in the car crash case. And we recognize that there was indeed a prior time at which Randy might have taken steps to avoid becoming distracted by a later car accident. But while Randy is focused on getting to the store on his way home, there is no reason for him to think that placing additional reminders around him or on his iPhone would significantly raise the likelihood of remembering the milk stop. Yes, accidents sometimes happen, but rarely. The chances that he would witness an accident that erases or severely interferes with the normal operation of his intermediate-term memory are extremely small. This is what explains the fact that in the car-crash case Randy is not morally responsible for his failure to stop for milk. (Notice here the way in which the car-crash version of Randy's case differs from the original case. In the original case, Randy's failure to remember to stop for milk is the result of his having become distracted by deep philosophical thoughts. At the time he started thinking about work while driving, Randy knew that philosophical thinking might well end up occupying his whole mind, and therefore that he would significantly raise his chances of not forgetting by either keeping his thoughts from wandering into philosophical territory or by setting a reminder on his cell phone.)

There are other ways in which events might rob an agent of the opportunity to do something that, as the agent reasonably believes, will raise the

likelihood of avoiding a later unwitting omission. An agent might be stricken by amnesia or paralysis, might be overcome by fear, depression, or some sort of compulsive disorder. An agent might be misled by evidence from a usually reliable source whose reliability she has no good reason to question. Imagine, for instance, that Alessandra, having entered her children's school, runs into a friend who tells her that she saw Alessandra's husband playing with the family dog in the parking lot. Alessandra has no good reason to disbelieve her friend, but, as it turns out, her husband wasn't playing with the family dog, but with a neighbor's dog (of a similar size and breed). Is Alessandra morally responsible for omitting to go back to the car? No, for she has every reason to believe that her husband is taking good care of the dog. Or imagine that ophidiophobic Alessandra meets a snake in the school hallway and, as a result, forgets about her dog. Or that Alessandra suffers a ministroke that deprives her of the ability to keep her mind focused on her dog's needs. In all these cases, Alessandra is off the moral hook, for she does not satisfy one or more of the conditions necessary for moral responsibility for her later omission.

The flip side of this objection is that on the Opportunity Tracing view almost no one is responsible for unwitting omissions, since in a large number of such cases there is no prior moment at which one was aware of the risks of the later omission coming to pass. After all, as we saw, requiring awareness is what drives some to skepticism.

But here we are inclined to the view that while people are often not aware of the risk of a later omission with the uniquely identifying, or even the specific, features of the unwitting omission that actually takes place, this sort of specificity is not required for responsibility to trace back. So while in the cases we have discussed we have indeed supposed that both Alessandra and Randy are aware of fairly specific risks (concerning the dog and the milk) ahead of time, in other cases less specificity might be required. (In this, we are in agreement with Fischer and Tognazzini [2009].) So, to take a case of FitzPatrick's (based on *It's a Wonderful Life*), if an agent, call him "Potter," had an opportunity to prevent his later unwitting failure to help his fellow townspeople, that opportunity needn't have been one that included his foresight that he was in danger of causing particular individuals' bankruptcies. Rather, his opportunity might involve simple awareness that failure to take steps to check himself now runs a serious risk of disregarding others' interests later. Thus, we take it that the view is not so restrictive that skepticism of even a significantly qualified kind results.

Both of the previous objections raise the interesting question of the conditions under which awareness can be attributed and still suffice for

providing opportunity. Does one have to be paying attention to the relevant facts? Can one be aware without rehearsing words to oneself? We believe that one can be aware in the relevant sense without doing so, and recognition of this point is crucial for determining the scope of what we are responsible for. But we acknowledge that there is much interesting work to be done to understand the different ways one can be aware of potential future risks.¹³

A third sort of objection is that while the basic account we give of control, and of moral responsibility, is correct, we have mistakenly required awareness for control, and for the possession of opportunity. Clarke (this volume) suggests that if Randy has the ability while driving to remember his earlier promise to stop for milk, then it is reasonable to expect him to *think* to stop for milk as he nears the store. According to Clarke, this ability is one element of what grounds Randy's basic, nonderivative moral responsibility for his unwitting omission to go to the store. On this view, as on the Below Standard view, what accounts for moral responsibility is what the agent is or is not able to do *at the time of the unwitting omission*: tracing by itself cannot do all the necessary theoretical work. One might think that there is a kind of control here, and that it does not itself require actual awareness. One has, in an important sense, the ability to think about stopping for the milk.

But there are two senses of "ability": general and specific. If Serena Williams is tied up, she retains the general ability to play tennis, but loses the specific ability to play tennis there and then. In one sense of "can" she can play; in another sense of "can," she can't. What matters for moral responsibility, surely, is specific ability, rather than general ability. If Harry is choking on some food but Voldemort has tied Hermione to her chair, she is not morally responsible for her failure to walk over to Harry and administer the Heimlich maneuver, even though she retains the general ability to do so. What absolves Hermione is that she has lost the *specific* ability to help Harry.

So let us ask whether Randy retains the *specific* ability to remember his promise to stop for milk *while his mind is occupied with philosophical thoughts relating to his book on omissions*. This is a difficult question, but there is a good case to be made for the claim that he has lost the relevant capacity (even if only temporarily). In the first place, remembering is not something that Randy *does*; it is, rather, something that happens to him. He can jog his memory, he can engage in mental calisthenics, he can go over recent events in his head: these are the sorts of mental activities he can perform

13. See Mole (2013, section 3.1) for an overview of the debate about the relationship between attention and consciousness.

voluntarily. But whether he remembers or not is not something within his direct control. He cannot simply will himself to remember. That is not how memory works. And while Randy is in the depths of intense philosophical focus, constructing hypothetical scenarios, thinking of ways of parrying potential counterexamples, concentrating on how best to articulate his own views and the arguments for them, he may not, in fact, be *specifically* capable of recalling his earlier promise, even though he hasn't lost the *general* ability to recall it. Thus, we would need more explanation of the grounds for claiming that Randy can retain control in the relevant sense without either awareness at the time, or awareness at an earlier time of the risk of a failure to be aware later.

On our view, genuine opportunities require awareness, so an unwitting omission, by definition, cannot be one that one has the opportunity to avoid *at the time of the omission*. Though it is often true that we think people are responsible when they *should* have been aware, though they are not actually aware, this is because their ignorance is culpable. And what makes it culpable, and what makes it the case that it is appropriate to say they should have been aware, is that at a previous time, they had the opportunity to be aware at the later time.

Finally, one might object that one cannot have control in a situation over something in a case in which one cannot exercise it; and in the case at hand, we are taking it that one has the relevant control at the time to which responsibility for the later omission traces back, namely, the time at which one is aware of risks and does nothing to minimize them. But Zimmerman (this volume) entertains the objection (without endorsing it) that one's failure to do anything at that time is not something one *can* do. It isn't an exercise of agency at all. Our answer to this objection is that because the Opportunity Tracing view is asymmetrical, one has all the control required for responsibility if one has the opportunity to do the right thing for the right reasons. One needn't have a kind of "dual control" in order to have all that one needs in the way of control for moral responsibility.

6.7. Conclusion

The commonsensical thesis that persons can be morally responsible for their unwitting omissions represents something of a philosophical nut. As we have argued, those who have tried to crack the nut, including basic theorists of various sorts (namely, Attributionists, Below-Standardists, and Underlying Self theorists) and tracing theorists (namely, Decision Tracers and Akrasia-or-Vice Tracers), face significant theoretical challenges. We have

tried to meet these challenges by offering a new tracing theory of moral responsibility for unwitting omissions that keeps all the advantages, while avoiding the disadvantages, of past proposals. On the Opportunity Tracing view that we endorse, and have attempted to defend, agents' moral responsibility for their unwitting omissions traces back to a time when they had the opportunity to do something that, as they reasonably believed, would significantly raise the probability of avoiding later omission of a certain kind. The opportunity to A, as we understand it, involves the witting interference-free ability to A, that is, the witting ability to A in circumstances that are sufficiently amenable to A-ing. On this view, we can accommodate the intuition that moral responsibility requires control, even though unwitting omissions are not under agents' control at the time they occur. And we can explain all of the intuitive judgments about particular cases that have been used to support alternative theories, of both the basic and tracing kinds.

References

- Agule, Craig K. 2016. "Resisting Tracing's Siren Song." *Journal of Ethics and Social Philosophy* 10: 1–24.
- Björnsson, Gunnar. Forthcoming. "Explaining (Away) the Epistemic Condition on Moral Responsibility." In *Responsibility: The Epistemic Condition*, edited by Philip Robichaud and Jan Willem Wieland (Oxford: Oxford University Press).
- Brink, David O. and Nelkin, Dana K. 2013. "Fairness and the Architecture of Responsibility." In *Oxford Studies in Agency and Responsibility*, edited by David Shoemaker (Oxford: Oxford University Press), pp. 284–313.
- Clarke, Randolph. 2014. *Omissions: Agency, Metaphysics, and Responsibility*. New York: Oxford University Press.
- Clarke, Randolph and Capes, Justin. "Incompatibilist (Nondeterministic) Theories of Free Will." *Stanford Encyclopedia of Philosophy* (Fall 2015 edition), edited by Edward N. Zalta. <http://plato.stanford.edu/archives/fall2015/entries/incompatibilism-theories/>.
- Fischer, John Martin and Ravizza, Mark. 1998. *Responsibility and Control: A Theory of Moral Responsibility*. Cambridge: Cambridge University Press.
- Fischer, John Martin and Tognazzini, Neil A. 2009. "The Truth about Tracing." *Noûs* 43: 531–56.
- FitzPatrick, William J. 2008. "Moral Responsibility and Normative Ignorance: Answering a New Skeptical Challenge." *Ethics* 118: 589–613.
- Ginet, Carl. 2000. "The Epistemic Requirements for Moral Responsibility." *Philosophical Perspectives: Action and Freedom* 14: 267–77.
- Khouri, Andrew C. 2012. "Responsibility, Tracing, and Consequences." *Canadian Journal of Philosophy* 42: 187–208.

- King, Matt. 2009. "The Problem with Negligence." *Social Theory and Practice* 35: 577-95.
- Levy, Neil. 2009. "Culpable Ignorance and Moral Responsibility: A Reply to FitzPatrick." *Ethics* 119: 729-41.
- McKenna, Michael. 2008. "Putting the Lie on the Control Condition for Moral Responsibility." *Philosophical Studies* 139: 29-37.
- Mole, Christopher. 2013. "Attention." *Stanford Encyclopedia of Philosophy* (Fall 2013 edition), edited by Edward N. Zalta. <<http://plato.stanford.edu/archives/fall2013/entries/attention/>.
- Nelkin, Dana Kay. 2011. *Making Sense of Freedom and Responsibility*. Oxford: Oxford University Press.
- Nelkin, Dana Kay. 2013. "Desert, Fairness, and Resentment." *Philosophical Explorations* 16: 117-32.
- Nelkin, Dana Kay. 2016. "Difficulty and Degrees of Moral Praiseworthiness and Blameworthiness." *Noûs* 50: 356-78.
- O'Connor, Timothy. 2016. "Free Will." *Stanford Encyclopedia of Philosophy* (Summer 2016 edition), edited by Edward N. Zalta. <http://plato.stanford.edu/archives/sum2016/entries/freewill/>.
- Rosen, Gideon. 2004. "Skepticism about Moral Responsibility." *Philosophical Perspectives* 18: 295-313.
- Shabo, Seth. 2015. "More Trouble with Tracing." *Erkenntnis* 80: 987-1011.
- Sher, George. 2009. *Who Knew? Responsibility without Awareness*. New York: Oxford University Press.
- Smith, Angela M. 2005. "Responsibility for Attitudes: Activity and Passivity in Mental Life." *Ethics* 115: 236-271.
- Smith, Angela M. 2008. "Control, Responsibility, and Moral Assessment." *Philosophical Studies* 138: 367-92.
- Smith, Holly. 1983. "Culpable Ignorance." *Philosophical Review* 92: 543-71.
- Vargas, Manuel. 2005. "The Trouble with Tracing." *Midwest Studies in Philosophy* 29: 269-91.
- Wallace, R. Jay. 1994. *Responsibility and the Moral Sentiments*. Cambridge, MA.: Harvard University Press.
- Watson, Gary. 2004. "Two Faces of Responsibility." In *Agency and Answerability: Selected Essays* (Oxford: Clarendon Press), pp. 13-32.
- Wolf, Susan. 1990. *Freedom within Reason*. New York: Oxford University Press.
- Zimmerman, Michael J. 1997. "Moral Responsibility and Ignorance." *Ethics* 107: 410-26.

PART III

Omissions, Moral Responsibility, and Alternative Possibilities