

RESPONSIBILITY AND VIGILANCE
(forthcoming in *Philosophical Studies*)

Samuel Murray, University of Notre Dame

ABSTRACT

My primary target in this paper is a puzzle that emerges from the conjunction of several seemingly innocent assumptions in action theory and the metaphysics of moral responsibility. The puzzle I have in mind is this. On one widely held account of moral responsibility, an agent is morally responsible only for those actions or outcomes over which that agent exercises control. Recently, however, some have cited cases where agents appear to be morally responsible without exercising any control. This leads some to abandon the control-based account of responsibility and replace it with an alternative account. It leads others to deny the intuition that agents are responsible in these troublesome cases. After outlining the account of moral responsibility I have in mind, I look at some of the arguments made against the viability of this theory. I show that there are conceptual resources for salvaging the control account, focusing in particular on the nature of vigilance. I also argue that there is empirical data that supports the control account so conceived.

Introduction

My primary target in this paper is a puzzle that emerges from the conjunction of several seemingly innocent assumptions in action theory and the metaphysics of moral responsibility. The puzzle I have in mind is this. On one widely held account of moral responsibility, an agent is morally responsible only for those actions or outcomes over which that agent exercises control. Recently, however, some have cited cases where agents appear to be morally responsible without exercising any control. This leads some to abandon the control-based account of responsibility and replace it with an alternative account. It leads others to deny the intuition that agents are responsible in these troublesome cases.

After outlining the account of moral responsibility I have in mind, I look at some of the arguments made against the viability of this theory. I show that there are conceptual resources for salvaging the control account. In particular, I focus on the nature

of vigilance. Vigilance is a disposition to become occurrently aware of morally or prudentially relevant considerations that constitute a sufficient reason to act or omit. I sketch the role that vigilance plays in our moral agency and indicate the way in which appeals to vigilance can expand the explanatory scope and plausibility of control-based accounts of responsibility. I also suggest that there is empirical data that supports my account of vigilance.

I. The Capacity Account of Moral Responsibility

Consider three variations on a story.

Josie, your infant child, is sitting in her high chair at dinner. During the meal, she spills some of her food onto the carpet.

Chuckie, your not-yet-house-trained puppy, is sitting at your feet during dinner. During the meal, she goes to the bathroom on the carpet.

Paul, your adult roommate, is eating dinner with you. During the meal, he throws his plate, full of food, directly at the wall.

Everyone agrees that Josie and Chuckie are not morally responsible for what they do.¹

That is, nobody would blame them for what they did. Everyone agrees that Paul *is*

¹ In this paper, I am concerned solely with moral responsibility. Any talk of ‘responsibility’ is shorthand for ‘moral responsibility’. Additionally, I am concerned solely with the accountability sense of responsibility, rather than the attributability or answerability senses (cf. Shoemaker 2015), where the accountability sense of responsibility picks out a range of attitudinal responses bound up in non-trivial ways to the Strawsonian reactive attitudes of resentment, indignation, and guilt (Strawson 1962). Finally, I examine only retrospective responsibility, or responsibility for what one has done.

(absent some excusing condition) responsible for what he does.² The natural explanation for the different assessments is that Josie is a child and Chuckie is an animal, while Paul is an adult. Adults are responsible for their behavior while infants and animals are not.

The reason behind the difference is that adults possess a wider range of competences and abilities that underwrite certain expectations we have about what people will do in certain situations. Infants and animals simply do not possess these competences and abilities. So, the reason that we treat Paul differently from Chuckie and Josie (in this context) is that Paul possesses the ability to recognize and behave in accordance with the expectations that people have of him. Chuckie and Josie have no comparable ability, so they are not responsible for what they do.

This explanation points to a particular general theory of responsibility, namely that agents are responsible in virtue of possessing and exercising certain agential capacities.³ The capacities in question are just those that enable adults to engage in forms of behavior that children and animals cannot. Let's call this theory, for ease of reference, the *capacity account of responsibility* (hereafter, Capacity Account), because the theory takes as central to the nature of responsibility the possession and exercise of particular agential capacities.

The Capacity Account specifies two broad classes of capacities the possession of which constitutes moral agency. The first class comprises those capacities needed for self-governance. These are the capacities under discussion in standard free will debates. The second class comprises those capacities needed for agential awareness. These are

² This is somewhat overstated. I mean that for anyone who holds that people are sometimes morally responsible for what they do, that person will agree that Paul is morally responsible for what he does.

³ Fischer and Ravizza (1998); van Inwagen (1983); Vargas (2013); Wallace (1994); Wolf (1990); Nelkin (2008); Raz (2011).

broadly epistemic capacities like perception, foresight, means-end reasoning, etc.

Possession of a suitable range of capacities across these two classes constitutes moral agency.

These capacities are central to moral agency (on the Capacity Account) because these capacities enable agents to recognize, appreciate, and respond to moral reasons. The epistemic capacities are needed to perceive reasons and appropriately weigh reasons against each other in deliberation. The self-governance capacities are needed to align one's behavior with one's view of the arrangement of moral reasons. The Capacity Account, then, implies that moral agency comprises a range of abilities that enable agents to appropriately self-govern themselves in the light of reasons.

At a high level of generality, the Capacity Account posits the following conditions for responsibility for actions and outcomes.

- (1) Agent is morally responsible for x (some action or outcome) when Agent possesses control with respect to x or with respect to some process p that normally and reliably issues in x , and;
- (2) Agent is morally responsible for x when Agent sees (or foresees) or should have seen (or should have foreseen) that either:
 - a. x 'ing is itself either morally right or wrong, or
 - b. x 'ing or bringing it about that p would likely lead to the occurrence of some morally undesirable or desirable state of affairs.

These conditions correspond to the two classes of capacities that jointly constitute moral agency. Moral agents are responsible for their conduct in virtue of fulfilling these conditions at the relevant time. Thus, the Capacity Account establishes symmetrical conditions on moral agency and responsibility. That is, the conditions on responsibility for actions (and responsibility for outcomes) fall out neatly from the account of moral agency.

II. *Arguments against the Capacity Account*

The Capacity Account follows nicely from certain intuitions about responsibility and has a wide explanatory scope with a small number of conditions. Recently, however, some have pointed out problems with the Capacity Account. These problems target either the plausibility of the Capacity Account or the scope of the theory.

Recall that the Capacity Account implies that human beings exhibit a particular form of moral agency constituted by agential capacities that fall into two distinct classes. A *plausibility argument* against the Capacity Account targets the theory of agency and uses these arguments to undermine the plausibility of the theory of responsibility. John Doris, for example, argues that the conditions on moral agency posited by the Capacity Account are such that agents rarely (if ever) fulfill such conditions.⁴ He claims that agents rarely exhibit self-control (of the sort specified by the Capacity Account) over their actions because situational factors of which the agent is unaware causally influence decision-making. Because one's decisions are subject to influence by a large number of factors that operate below the level of awareness, this precludes control over one's decisions and, by extension, behavior. Thus, Doris argues that if the Capacity Account were true, then we would rarely be responsible because we rarely fulfill the conditions on moral agency. But we are frequently responsible, so the Capacity Account (and its attendant theory of moral agency) must be false.

Scope arguments claim that the Capacity Account generates assessments of certain cases that conflict with our intuitions. In general, these cases describe some agent

⁴ Doris (2015: 43).

that appears to behave in some morally significant way despite lacking or failing to exercise the capacities that constitute moral agency.⁵ Here's one example, taken from Randy Clarke's work:

Milk. As I'm about to leave my office at the end of the workday, my wife calls to tell me we're out of milk. My regular route home takes me right by a grocery store, and I tell her I'll stop and buy some. Between my office and the store, I start to think about a paper I'm writing on omissions. I continue thinking about my work until I arrive home, where I realize that I've forgotten the milk.⁶

Let's name the individual in this story Joel. Joel is responsible for failing to pick up milk despite being unaware of this failure. Thus, at the relevant time, Joel is unaware that his behavior is likely to issue in an undesirable state of affairs. Hence, Joel fails to satisfy condition (2.b) of the Capacity Account. So, the Capacity Account cannot accommodate the intuition that Joel is responsible.

There are two ways for proponents of the Capacity Account to respond to these arguments that are initially unsatisfying. The first is to introduce a tracing principle. Traces allow proponents of the Capacity Account to connect some moment where an agent fails to fulfill the Capacity Accounts conditions on responsibility back to some prior moment where the agent fulfills the Capacity Account conditions on responsibility and behaves in a way that explains the later failure to exercise her capacities. So, for example, we can explain Joel's failure to get the milk in terms of some prior behavior that explains why Joel fails to get the milk.⁷ Similarly, proponents of the Capacity Account can note that even if people rarely possess control over and awareness of their decisions and actions, it is implausible to suppose that agents *never* possess control and

⁵ See Frankfurt (1998), Scanlon (1998, 2008), Smith (2005, 2008, 2015)

⁶ Clarke (2014: 164). Clarke does not cite this example in support of a scope argument.

⁷ For some attempts, see Ginet (2000), Fischer and Tognazzini (2009), and Timpe (2011).

awareness. Thus, while rare, agents exhibit enough control and awareness to anchor responsibility ascriptions in those moments where agents lack control and awareness.⁸

While the tracing mechanism might provide an explanation for certain failures, it is unclear that tracing will be ubiquitously applicable in the relevant cases. There appear to be cases where an agent is responsible for some behavior despite never fulfilling at some earlier time the Capacity Account conditions on responsibility in a way that explains an agent's responsibility when that agent lacks either control or awareness.⁹ In fact, *Milk* seems to be an example of just this sort! Reliance on tracing, especially in cases like *Milk*, seems misguided. Another reason for skepticism about the prospects of applying a tracing mechanism to explain responsibility in these cases is that tracing does not fit into our phenomenology of blame. That is, we find it appropriate to blame the agent in question in the absence of identifying a potential anchor for a trace.¹⁰ For these reasons, a successful response on behalf of the Capacity Account ought to specify conditions on responsibility that do not include some tracing principle (at least, proponents of the Capacity Account ought not to include a tracing mechanism to explain cases of this sort; there may be other principled reasons to include a tracing mechanism that warrant the inclusion of this mechanism in one's theory).

Second, proponents of the Capacity Account might say that all of these arguments are easily rebutted by the inclusion of a normative element in condition (2.b). Thus, Joel is responsible because he *should* have been aware of needing to get the milk. The initial

⁸ Though his motivations were different, van Inwagen (1989) raised similar worries about the scarcity of control that agents possess. The suggestion I countenance above is the one he suggests. Even if the scope of controlled action is relatively small, we can use tracing mechanisms to maintain the relatively wide scope of moral responsibility.

⁹ Cf. Sher (2009: 83).

¹⁰ Doris (2015: 155). See also Graham (2014: 399).

formulation of the conditions on responsibility for the Capacity Account appears built to respond to objections of just this kind. So what's all the fuss about?

The problem with normative responsibility conditions is that these do not establish any clear link between the agent and her actions. To see this, consider a normal case of culpable wrongdoing. When an agent fulfills the responsibility conditions in this normal case, the moral qualities of the act figure into the set of considerations to which the agent is sensitive. When the agent acts wrongly in light of these considerations it is possible to draw a direct line from the negative moral qualities of the action to some moral feature of the agent. When, however, we invoke some normative condition to explain responsibility, this direct line disappears.

The *Milk* case makes this point more concrete. When Joel gets home, his partner reprimands him for not getting the milk. When Joel responds that he didn't realize what he was doing, we can imagine his partner saying: "Well you *should* have known." But what does Joel's partner mean here? Joel does not realize that he is failing to do something and certainly doesn't realize that he is violating some norm about paying attention to what he is doing. Thus, the partner appears to explain Joel's responsibility for failing to get the milk by appealing to some failure to discharge his epistemic duties. This, however, is just to explain responsibility for a failure by appealing to another failure. The explanation does not appear to provide any means for drawing a line of attribution from the wrongness of Joel's behavior to any positive feature of Joel himself.¹¹ So, the inclusion of some normative component within the responsibility conditions does not obviously answer any of the arguments offered against the Capacity Account.

¹¹ Sher (2009: 26-7, 74).

Some philosophers use plausibility or scope arguments against the Capacity Account to motivate adopting some form of skepticism¹² or revisionism¹³ about the nature of moral responsibility. I will not summarize these arguments here; instead, I propose to offer a response on behalf of the Capacity Account. In general, I claim that the Capacity Account has the resources to respond to the plausibility and scope arguments normally leveled against them. This, in turn, will preserve the isomorphism between responsibility and moral agency. I show that the Capacity Account can tell a plausible story about why agents in *Milk*-style cases are responsible for their behavior. To do this, I take up two questions. First, what sort of normative account explains responsibility in cases where agents seem to lack the capacities constitutive of moral agency? Second, what does it mean to possess the relevant suite of abilities connected to this normative account?

III. *Vigilance: The Conceptual Dimension*

Let me start this section by proposing a definition. Let some behavior *B* of an agent *S* be an *unwitting omission* when: (a) *B* is wrong or right in virtue of reason (or a set of reasons) *r*, and; (2) *S* *B*'s or fails to *B* without being aware of *r* under any morally relevant description.¹⁴ I introduce this term here because the arguments mentioned against the Capacity Account often invoke cases of unwitting omissions. This is to be

¹² Zimmerman (2008) and Rosen (2003, 2004).

¹³ Smith (2005, 2008, 2015) and Scanlon (1998, 2008). This form of revisionism is distinct from the revisionism of Vargas (2013) and Nichols (2015). Vargas and Nichols are not revisionists about the centrality of capacities to moral responsibility.

¹⁴ Here, the reasons are normatively connected to behavior. That is, in any context there will be normative reasons that speak in favor of certain forms of behavior and speak against others. Thus, when an agent unwittingly omits, the agent behaves in some way, but the reasons that speak against behaving in that way do not figure into the set of considerations to which the agent is actually sensitive. This account is designed to remain neutral between different theories of reasons. Internalists and externalists alike ought to be able to fill in the above proposal without any significant changes to the argument that I offer.

expected. After all, unwitting omissions by definition involve some failure to behave appropriately in light of some lack of awareness of the omitting agent. Alternatively, we could characterize an unwitting omission as any instance where an agent behaves in some morally significant way, but the moral significance of the behavior plays no causal or explanatory role in the agent's so behaving.¹⁵

The lack of awareness constitutive of unwitting omissions seems to suggest that the unwitting agent in question lacks certain self-control capacities. In other words, awareness appears to be a necessary condition for control, so without awareness one also *lacks* the capacities that fall within the self-control class.¹⁶ With that in place, we can begin to look at the response on behalf of the Capacity Account. I propose to start with the normative dimension of this problem. That is, in virtue of what are agents culpable for unwitting omissions?

The normative basis for moral responsibility for unwitting omissions may be plausibly taken to involve some obligation to maintain a certain level of vigilance. The obligation to maintain vigilance has two components. The first is the property or ability that the obligation picks out. This, on my account, is vigilance. The vigilance of an agent consists of a disposition to become occurrently aware of morally or prudentially relevant

¹⁵ One consequence of this characterization is that it entails that no morally neutral behavior could be characterized as an unwitting omission. For example, on this account, if someone steps into the shower with socks on this would not count as an unwitting omission (assuming there is nothing morally significant about getting into the shower with your socks on). Someone might take issue with this consequence and thus take issue with the characterization itself. Randy Clarke (2014: 94) suggests that one might unwittingly omit in the context of a game. This dispute feels somewhat verbal, though we could (I think) characterize unwitting omissions as a species of unintentional omission, where unwitting omissions are those unintentional omissions that are morally significant.

¹⁶ Even if responsibility theorists maintain a distinction between self-control and epistemic agential capacities, there is still a tendency to include some distinctive epistemic component in the self-control capacities, as in Vargas (2013: 201-202) and Ekstrom (2000: 67). For this reason, lack of awareness of some salient aspect of one's moral environment might suffice to count an agent as lacking relevant self-control capacities.

considerations that constitute a sufficient reason to act or omit.¹⁷ The exercise of this disposition constitutes the agent's attunement to moral or prudential reasons to act. When those reasons constitute an obligation to act or omit, failures to act or omit is culpable. That is, vigilance enables an agent to increase or intensify one's sensitivity to available moral reasons.¹⁸

The second component of the obligation is the level of vigilance that one is required to maintain. Levels of vigilance are measured in terms of how much effort one devotes to being aware of and attuned to one's environment and acting in light of that awareness. Thus, the obligation makes demands on how much effort one ought to devote to paying attention. For example, the demand for vigilance with respect to driving is relatively low when one is driving on a dry, straight, wide stretch of empty highway in western Wyoming. The demand for vigilance increases when one is driving on a narrow, crowded city street during the night rush.

What sets the level of vigilance in a context is likely a function of one's broader normative commitments. So, utilitarians, virtue theorists, and deontologists will likely have different explanations for shifting demands for vigilance. To see one way in which broader normative commitments might settle the level of vigilance demanded in a

¹⁷ Pereboom (forthcoming: 185-86) also explains responsibility for unintentional omissions in terms of failures of vigilance. His conception of vigilance, however, is much narrower than the one in this paper. Pereboom describes vigilance as: "a persisting attunement to protect, which features, among other things, a standing disposition to respond to danger, triggered by indications of danger in the environment" (185). Vigilance, as I describe it, includes this disposition but also includes a range of other dispositions that have no connection to awareness of needs to protect (of course, there could be some naturalistic explanation of the development of vigilance that posits Pereboom's disposition as a course-grained precursor to the more variegated and complex vigilance capacity that I countenance). Given Pereboom's narrow conception of vigilance, he expresses skepticism that vigilance can furnish a general explanation of responsibility for unintentional omissions (186). Because I offer a thicker conception of vigilance, I take it that my vigilance account *can* furnish such a general explanation.

¹⁸ Following Fischer and Ravizza's (1998: 69) discussion, this sensitivity comes in two dimensions: receptivity and reactivity. Vigilance is connected to both, though the latter dimension is unimportant for the purposes of this paper.

context, consider what a classical utilitarian might say. Roughly, the classical utilitarian will say that the level of vigilance demanded in a context is a function of whatever actions will result in the maximization of relevant goods.

While the demanded level of vigilance might be a function solely of some broader normative theory, the level might also be a function of more localized community practices. On this picture, the level of vigilance might be set by the expectations of some relevant subset of a particular community. These members of the community have certain reasonable expectations about how people will behave in certain domains and these expectations set the standard of vigilance that one ought to meet in those domains.¹⁹

The level of vigilance that one ought to maintain does substantial work in setting the normative basis for responsibility for unwitting omissions. Despite this theoretical load, it is difficult to offer a more substantive account without first adopting some controversial, higher-order normative commitments. Insofar as the vigilance account that I offer here aims to be relatively modular, I withhold offering a more substantive account of the level of vigilance that one ought to maintain.

A substantive account of the *content* of the norms of vigilance likely requires making more substantive higher-order normative commitments. However, more can be said about the *function* of the norms of vigilance while retaining the modularity of the account.

There are two important functional features of the norms of vigilance. The first is that the norms of vigilance function as *derivative* norms. These norms are derivative because they derive their status from primary norms. As I am using the distinction, a

¹⁹ Cf. Vargas (2013: 214).

primary norm is such that one allows that norm to play a role in one's cognitive economy simply in virtue of understanding the content of the norm. Thus, a norm is primary if, upon understanding the content of the norm, one needs no further reason to adopt or submit to that norm. A derivative norm, on the other hand, is such that one allows that norm to play a role in one's cognitive economy only if, upon understanding the content of the norm, one also has some further reason to adopt the norm. The norms of vigilance are derivative because one only accepts the demands of these norms in virtue of accepting the demands of other norms or adopting other commitments. No one would, I submit, accept the demand for vigilance without some background web of normative commitments that supported accepting such a demand.²⁰

Describing this dimension of the account requires a longer story. Roughly, as moral agents we seem committed to engaging in various activities that require us to allocate attention in efficient ways. For instance, we participate in meaningful relationships that sometimes require us to make and fulfill promises. And our commitment to living up to these relationships (by fulfilling our promises) does seem to require living up to some demand that is made on our attention. Thus, we seem to accept demands for vigilance in virtue of wider commitments that we have as moral agents. And without these wider commitments, there does not seem to be much reason to take on board a demand for vigilance. In this way, the norms of vigilance are derivative.

The second important functional feature of the norms of vigilance is that the norms are context-sensitive. The demand for vigilance is situationally indexed to account for a variety of environmental factors that might make it more difficult (in some sense) to display a certain level of vigilance. In particular, the demand is sensitive to the level of

²⁰ Cf. Raz (2011: 237-39) and Vargas (2013: 249; ms.).

cognitive effort needed to maintain a level of vigilance. The account needs this context-sensitive component to account for the fact that the same vigilance capacity might be more or less personally available to the agent based on features of the agent's environment.²¹ This feature of the account adds two things: (1) the same norm might make different demands on different individuals in the same context, and; (2) the same norm will make different demands on the same individual cross-situationally. These additions are needed to account for intuitions about certain cases. For instance, with this added dimension, we can make sense of why we excuse a foreigner for making an improper turn at a red light on a busy city street where we would not excuse a native for the same action. And, we might excuse someone for committing a social *faux pas* at a dinner party with royalty when we would not excuse them for committing a *faux pas* at a dinner party with friends and family. In general, the context-sensitivity of the demand for vigilance is a function of the stress that the environment puts on the vigilance mechanism. Given a certain threshold of stress, the demand is weakened to account for the additional effort needed to display appropriate levels of vigilance.

The normative account of responsibility for unwitting omissions identifies a relatively specific feature of an agent, vigilance, which I defined above as a disposition to become occurrently aware of morally salient features of the environment. The exercise of this disposition constitutes the agent's attunement to the moral environment. Having this agential capacity warrants the application of a cognitive standard to the agent across a wide range of circumstances (namely, any circumstance in which the agent can exercise and calibrate her level of vigilance). Agents are culpable for falling short of this cognitive

²¹ Note that this formulation remains neutral on whether situational factors themselves structure agential capacities or whether situational factors merely affect the operability of a particular capacity.

standard because they are obligated to maintain a certain level of vigilance over those capacities that the standard measures. Thus, the cognitive standard *measures* or *tracks* the level of vigilance that the agent manifests and *applies* to the agent in virtue of capacities that the agent possesses.

The norms of vigilance, on this account, are sensitive to two features of the vigilance capacity. On the one hand, the norms of vigilance correctly apply to some agent just because that agent possesses the appropriate disposition for vigilance. Insofar as vigilance is partially constitutive of moral agency, it follows that the norms of vigilance will only make demands on those agents that are capable of fulfilling the demands of those norms (this, in turn, follows from adopting the Capacity Account). On the other hand, the specific demands that the norms of vigilance make are a function of the degree to which one can fully exercise one's vigilance. It might be more difficult in certain situations to achieve a certain level of attunement to the moral environment. Thus, the norms of vigilance reflect two facts about the vigilance capacity: 1) vigilance is a capacity like any other that some creatures possess and others do not, and; 2) vigilance is a flexible capacity that requires more or less effort to appropriately exercise across a wide range of scenarios.

In sum, there are three components to this normative dimension of vigilance. First, moral agents possess some disposition (vigilance) the exercise of which affords them some degree of attunement to normatively salient features of the environment. Second, certain derivative moral norms specify a cognitive standard that determines what level of vigilance moral agents ought to display within a context. Third, agents are

justifiably held to the demands of the norms of vigilance in virtue of having the ability to be vigilant.

Hence, the vigilance account explains culpability for unwitting omissions by appealing to the agent's display of substandard awareness. One can see this in the *Milk* case. Joel is responsible for failing to bring home milk because he failed to exhibit the appropriate level of vigilance at the relevant time (and, we suppose, the circumstances were not such that the demand for vigilance made on Joel was unreasonable or unfair). This gives the vigilance account a certain theoretical advantage over revisionist positions. For instance, Angela Smith attempts to explain all responsibility for unwitting omissions in terms of a lack of due care.²² While this framework might explain some instances of responsibility for unwitting omissions, it seems to fail in other cases. After all, in *Milk* Joel simply seems to make a mistake.²³ He has all the right attitudes toward fulfilling promises and showing respect for his partner. Yet he still fails to get milk. When we describe the case as a mistake the vigilance account can still furnish an adequate explanation. Explaining responsibility in terms of vigilance, rather than due care, seems more appropriate.

Note that as I have explained it here, the vigilance account functions as an explanatory framework for intuitive judgments about the culpability of unwitting wrongdoers. That is, we intuitively judge that sometimes agents are culpable for their unwitting omissions. As I noted above, explaining these judgments is problematic on a standard construal of the conditions for moral responsibility on a Capacity Account. With the vigilance account, the Capacity Account has an explanation for these judgments on

²² See Smith (2005).

²³ Cf. Amaya (2013) and Amaya and Doris (2014).

hand. This account cannot, as described here, *establish* that an individual is culpable for an unwitting omission. Connecting the vigilance account back to some wider theory about the basis of culpability requires separate treatment.

IV. Vigilance: The Empirical Dimension

In the previous section, I proposed that the normative basis for moral responsibility for unwitting omissions involves the obligation to maintain a certain level of vigilance. This requires that agents possess some form of control over the levels of vigilance that they display (otherwise, the maintenance of vigilance would be something outside of the scope of one's agential capacities). In short, my proposal seeks to add and describe an additional capacity within the Capacity Account picture, one that accounts for vigilance.

In this section, I want to look at recent work on the neuroscience of cognitive control. This is important because there are a number of empirical claims that the normative dimension of the vigilance account implies. For one, the normative dimension of the vigilance account implies that we realize a certain form of agency that has a certain underlying architecture (i.e., one that supports the ability to be vigilant). Because this robust empirical component falls out of the normative account, we need to determine whether the normative account has any empirical traction.²⁴

To see why this is a worry, consider an analogy. Imagine a computer scientist that wants to build a device that can perform functions *f*, *g*, and *h*. The computer scientist designs a range of software to perform these functions. Now suppose that the computer scientist takes his idea and his software to an electrical engineer. The electrical engineer

²⁴ In this, I follow a suggestion made in Doris and Stich (2007). Normative theories often contain empirical elements or make strong implications about the way the world is. As such, normative theorizing ought to be constrained by empirical considerations to *some* degree.

could say that there is no way to build a machine that can run the computer scientist's software. That is, the electrical engineer might say that we do not possess the physical resources to craft something that can perform the desired functions. Something similar, I think, can be said in the case of moral agency. The normative dimension of the vigilance account implies a certain conception of moral agency. One reason to go to the neuroscience is to figure out whether we have reason to believe that any biological organism possesses the physical make-up to perform the tasks demanded by the vigilance account.²⁵

One long-standing assumption of neurologically grounded theories of behavior is that agents possess a disposition to behave in ways that will optimize the expenditure of resources with the expected benefits that accrue from such expenditures.²⁶ One way of minimizing resource expenditures is to limit the allocation of control across a variety of mechanisms. A relatively basic or low-level notion of control is given by the action of those mechanisms responsible for configuring behavior in such a way that maximizes reward attainment. Control, thus understood, is cognitively 'expensive' or resource demanding. So agents benefit from activating control processes as little as possible.

A recent model of cognitive behavior, articulated by Jonathan Cohen, provides a framework for understanding the allocation of cognitive demand. The model, the 'expected value of control' (EVC), predicts that allocation of control is based on the expected payoff from engaging a control process and the cost of achieving a payoff, where cost is analyzed (both psychologically and neurologically) in terms of exercising

²⁵ This proposal does not suggest reducing moral agency to neurobiology. I consider and respond to this concern in the next section.

²⁶ Kool et al. 2010; Botvinick 2007.

this cognitive effort.²⁷ EVC thus explains the integration of two sources of value-information: 1) expected payoff from activating (or signaling) a control process, and; 2) cost of engaging control mechanisms.

The EVC explains the way in which agents directly control their cognitive abilities by allocating cognitive demands across cognitive processing structures that exact differential cognitive loads. Important task representations are allocated to cues in ways that involve high cognitive load while less important representations are allocated in ways that involve low cognitive load. The EVC explains how one does or does not maintain a state of vigilance and thus accounts for the way in which one can exercise control (or fail to exercise control) over the level of vigilance that one displays.

Specifically, as mentioned above, important task representations (e.g., remembering anniversary dates, promises to get milk, or the location of car keys, etc.) can be encoded in ‘expensive’ ways that increase the likelihood of recall at the relevant time.

Mechanisms specified by this framework realize the ability to exercise vigilance.

The EVC provides not only a way of explaining the allocation of cognitive resources; it also identifies a single system of mechanisms responsible for processing the inputs in the EVC model. According to Cohen, the dorsal anterior cingulate cortex (dACC) integrates the various value inputs and determines where and how much control should be allocated to different cognitive tasks. The dACC interacts with other valuative and regulative systems to activate specific task units with control signals. In particular, the dACC regulates lower-level information processing mechanisms and monitors changing information about circumstances and signals that indicate whether certain

²⁷ Botvinick and Cohen 2014; Shenhav, Botvinick, and Cohen 2013.

control signals ought to be readjusted.²⁸ The dACC also communicates with the lateral prefrontal cortex (IPFC), which latter implements ‘decisions’ made by the dACC to allocate control.²⁹ The dACC functions as a mechanism that realizes the agential ability to calibrate degrees of vigilance.

This interaction of psycho-physical processes provides a mechanism through which agents can exercise and calibrate degrees of vigilance. Vigilance, here, picks out a virtue of agents that is realized in these underlying neural structures and activities and provides for the awareness of the environment that the agent possesses. Agents can exercise and calibrate degrees of vigilance in three ways: (1) gathering more information to shift reward expectation; (2) searching for alternative sources of reward, or; (3) valuing long-term over short-term rewards. In this way, a crucial aspect of vigilance is explicable in terms of the EVC.

The functional account of vigilance does not settle various normative questions about how much vigilance is required and under what circumstances. What it does provide, however, is a conceptually and empirically adequate realizer for fulfilling some independently specified normative account of the duty to be vigilant. The normative account correlates with capacities that can be cashed out in empirically tractable ways. The mechanisms identified under the EVC model accounts for the moral valence of a variety of cognitive, perceptual, and memory failures.

Notice that this functional account of vigilance implies intuitive excusing conditions and explains why these conditions excuse. Normally, duress counts as an

²⁸ Empirical confirmation in Bogacz et al. 2010; Forstmann et al. 2008, 2010; Ivanoff et al. 2008; Molder et al. 2012; van Maanen et al. 2011.

²⁹ Empirical confirmation in Aron et al. 2007; Aron and Poldrack 2006; Cavanagh et al. 2011; Jahafri 2011; Wiecki and Frank 2013; King et al. 2010; Kerns 2006. ‘Decisions’, here, is shorthand for computations performed by the dACC that issue in signals to implement relevant task representations.

excuse for morally undesirable behavior. Within the framework of the EVC, this is understandable. Given certain stressors in the environment, one's vigilance mechanism might be overwhelmed and thus fail to operate. For example, suppose that instead of thinking about a paper, Joel was thinking about a phone call that he received informing him that his father had died earlier that day. Under these circumstances, we would likely excuse Joel because his capacity to recognize and respond to his surroundings is swamped. Thinking about a paper does not overwhelm the vigilance mechanism in the same way, so we do not accept thinking about a paper as an excuse for failing to realize that one needs to go to the store. Again, the threshold that sets the standard for when the mechanism is overwhelmed likely results from some combination of objective facts about the cognitive effort needed to do what one ought *and* facts about what the community expects of similarly situated agents.

The functional account of vigilance actually implies an empirically verifiable hypothesis. Given that the model predicts that excuses for unwitting omissions ought to track the level of stress placed on an agent's vigilance mechanism, we can test whether our folk psychological categories of excuses matches up with this prediction. This empirical work has not yet been done, though I suspect that there will be significant overlap between intuitions about whether an excusing condition obtains (in the case of an unwitting omission) and the level of stress that an environment places on an agent's vigilance mechanism. There will likely be some slippage, i.e., cases where intuitions suggest excusing where the level of stress on the system is actually mild and cases where intuitions suggest withholding excuses where the level of stress is actually severe. However, absent some drastic divergence of the neuroscience from our folk

psychological categories, the EVC seems poised to vindicate this feature of our ordinary practices.

The development of the vigilance model provides a framework that explains why agents are morally responsible for their unwitting omissions. In these cases, agents are responsible because they *should* be aware of what they are doing but fail to be aware. The normative account of vigilance fills out what the relevant sense of ‘should’ is and posits a distinctive capacity, vigilance, the operation of which enables the agent to fulfill the duty to be aware. The empirical account of vigilance furnishes a mechanistic sketch of this normative capacity and specifies particular neural realizers for the relevant capacity. Given that agents possess this ability, we can see why the normative epistemic condition picks out a direct relationship between the agent and her behavior in cases of unwitting omission. When, for instance, Joel forgets the milk, there is a direct line between his failure to get milk and his vigilance capacity.

The addition of the vigilance account makes clear the response on behalf of the Capacity Account to the scope argument. The problem that the scope argument raises is that the Capacity Account generates assessments of particular cases that do not track our intuitive judgments about responsibility in an entire class of cases (namely wrongful unwitting omissions). The vigilance account provides resources to take this objection head-on. It does this by expanding the scope of the Capacity Account and provides principled reasons for accepting the machinery that accounts for this expansion. That is, the vigilance account countenances a specific feature of an agent (vigilance) that accounts for the culpability an agent bears for her unwitting omission. While the account appeals

to an unexercised capacity in these cases, the neuroscientific evidence provides support for the claim that agents possess this capacity even when they fail to exercise it.³⁰

The vigilance model also provides the resources to respond to the plausibility argument against the Capacity Account. Recall that the plausibility argument targeted the theory of agency that the Capacity Account implies. Two responses to this criticism emerge from the foregoing account. First, one could directly deny the claim that agents realize the form of agency specified by the Capacity Account. Given that the neurobiological mechanisms are ordinary features of the human being, it seems implausible to say that agents rarely utilize these mechanisms to configure their behavior in controlled ways. Given that the mechanisms within the EVC are directly sensitive to valuative systems in the brain, it also seems odd to say that an agent's view of reasons for acting (plausibly construed, at the neural level, as information about expected reward) plays an insignificant role in the production of action.

This response might be unsatisfying to some, but there is another response available in light of the vigilance account. The context-sensitivity of the norms of vigilance could address concerns raised by Doris and others. Because the norms of vigilance are sensitive to the level of stress placed on the vigilance mechanisms by features of the environment, it is open to proponents of the Capacity Account to say that demands shift to account for 'noise' that diminishes the availability of the vigilance mechanism and might make the exercise of control harder in a particular context. One could interpret the causal influence that situational factors exert on decision-making as noise that diminishes the demand for vigilance. The social psychology research that Doris

³⁰ This response is not likely to convince anyone skeptical of the idea of an unexercised capacity. Despite this, the response does show that the account has something principled to say for itself.

utilizes to fuel his plausibility argument might pick out a relatively isolated set of cases where our folk psychological categories fail to properly track the level of stress placed on an agent's vigilance mechanism. Thus, instead of claiming (as Doris does) that we rarely exhibit the relevant form of control required for responsibility, we can offer a slightly revisionary account of control and make the case that our folk psychological categories fail to appreciate just how sensitive the mechanisms that mediate this control are to environmental noise.

More could be said, but the general picture is that the vigilance account offers some plausible routes to answer these arguments against the Capacity Account. This is because: (a) the vigilance account expands the explanatory scope of the standard construal of the Capacity Account, and; (b) the vigilance account contains a context-sensitive element that accounts for the sensitivity that responsibility-relevant capacities display toward features of the environment.

V. Objection 1: Reductionism and Agency

In the next two sections, I want to briefly respond to two concerns with the vigilance account. The first objection concerns the relationship between the normative and empirical dimensions of the vigilance account and whether the normative dimension simply reduces to the empirical dimension. The second objection concerns whether the vigilance account provides the Capacity Account with an explanation of original (direct) or derivative moral responsibility for unwitting omissions. I take up the first objection in this section and the second objection in the next.

One concern with the relationship between the empirical and normative dimension of vigilance is that the above account appears to reduce the normative (and thus some crucial aspect of our agency) to the empirical. A related concern is that this reduction invites a form of skepticism. After all, if vigilance just is neurobiological functioning, then what contribution does the *agent* herself make (and in what way can the *agent* herself calibrate and display varying levels of vigilance)?

This is an extremely difficult problem and will likely arise for any theory of responsibility, agency, or free will that incorporates some robust empirical component. By way of response, recall first that the purpose of including the empirical dimension of the vigilance account was twofold. First, the normative dimension of the vigilance account implies a number of empirical theses. Thus, the plausibility of the normative dimension of the account *demand*s that we sketch at least some of the empirical implications of that dimension. Second, insofar as the empirical dimension confirms some of the empirical implications of the normative dimension, we have good evidence for the accuracy of the normative story. In particular, the empirical dimension supports the claim that agents possess a capacity for vigilance even when that capacity sometimes goes unexercised. Without the support of the empirical dimension, the normative dimension would seem hollow and speculative.

So there is good reason to include the empirical dimension of the account and thus good reason to explain away the worries about reductionism. The two dimensions of the above account can be thought of as descriptions of phenomena at various levels. In particular, the normative dimension is a description at a *psychological* or *agential* level, where our description includes terms like ‘agent’, ‘ability’, and ‘norm’. I do not mean to

imply, by this, that all psychological or agential descriptions are normative; rather, I mean that in explicating the normative content of the vigilance account, I took as basic particular psychological or agential concepts from which we can draw substantive conclusions. The empirical dimension is a description at a *neurological* level, where our description includes a number of mechanistic sketches that correlate to the activities at the psychological/agential level.

The neural mechanisms, in my view, *realize* the capacity for vigilance and *mediate* the agent's control of her unwitting omissions. Additionally, these mechanisms function as *neural correlates* for the exercise (and failure to exercise) the capacity for vigilance. The language of realization and mediation is crucial here. Neither term implies or entails any kind of reduction to more fundamental entities.

One might think that the explanatory relevance of the empirical dimension to the normative dimension entails some kind of reduction, insofar as we can reduce the phenomenon of vigilance to the activity of its underlying neural mechanisms. This worry, however, does not apply to the vigilance account. The empirical dimension of the account specifies neural mechanisms that explain the phenomenon of allocation of cognitive demand. These mechanisms do not explain the *ability* posited by the normative dimension of the account; rather, the neural mechanisms are just neural correlates of the capacity for vigilance. Merely specifying neural correlates, however, is insufficient for causal explanation.³¹ Thus, the empirical dimension does not provide a causal explanation for the normative dimension of vigilance (in particular, the capacity for vigilance specified by the normative dimension of the account); rather, it provides a

³¹ Craver (2007: 60).

mechanistic sketch for a particular cognitive phenomenon, from which sketch we can plausibly infer the presence of the very capacity specified by the normative account.

This response, however, does not account for the agent. Where does the agent fit in given the empirical dimension of the account? While this is an important lacuna in the vigilance account presented here, it is not clear that the objection is relevant. As I presented it, the vigilance account is neutral between substantive theories of agency and action. Thus, in my view, one can plug in one's preferred account of agency and theory of action and adjust the vigilance account accordingly. Because of this, it is not clear that the vigilance account provides some novel response (or is required to provide a novel response) to the skeptic over and above the responses already available from higher-level disputes.

VI. *Objection 2: Derivative and Original Responsibility*

The second objection concerns the kind of explanation offered by the vigilance account for unwitting omissions. This framework appears to offer explanations of *derivative* responsibility for unwitting omissions (where derivative responsibility for some x depends on some prior moment of original responsibility for some y that is suitably and foreseeably connected to x). That is, the responsibility for an unwitting omission is derivative on one's responsibility for a failure of vigilance. The problem with this is that all of the problems with tracing (which is a necessary component of explanations of derivative responsibility) mentioned in sec. 2 flood back in if we explain responsibility for unwitting omissions only derivatively. In addition to this, if the vigilance account explains responsibility for unwitting omissions merely derivatively, then it is not clear

that the vigilance account offers a distinctive explanatory framework for responsibility for unwitting omissions.

There are two responses to this line of objection. First, even if the vigilance account does rely on tracing mechanisms to explain responsibility for unwitting omissions, the account goes beyond its predecessors by specifying a tracing anchor that is sufficient for grounding responsibility ascriptions. The tracing anchor in this case is some moment where the agent fails to exercise vigilance. Specifically, it would be some moment where the agent fails to properly exhibit a certain degree of awareness or attunement to the moral environment. Anchoring responsibility for future unwitting omissions in these failures avoids some regress problems raised against other tracing explanations of unwitting omissions. This means that tracing is, on the vigilance account, less theoretically problematic than we initially supposed (and less problematic than alternative tracing schemas).

Second, it is not clear to me that the vigilance account does furnish explanations only derivatively. To see this, consider a recent proposal of sufficient conditions on original responsibility for unwitting omissions given by Randolph Clarke. Clarke suggests that there are three conjointly sufficient conditions:

Provided that the agent has the capacities that make her a morally responsible agent, she is blameworthy for such an omission if she is free in failing to [do] the thing in question and if her lack of awareness of her obligation to do it—and of the fact that she isn't doing it—falls below a

cognitive standard that applies to her, given her cognitive and volitional abilities.³²

Clarke's proposal contains three conditions that conjointly suffice for an agent to count as non-derivatively blameworthy:

- (i) Agent possesses capacities necessary to count as a moral agent
- (ii) Agent is free in omitting
- (iii) Agent's awareness falls below an applicable cognitive standard³³

Explanations generated by the vigilance account conform to this standard. Consider Joel again. At the time of his omission, the vigilance account claims that Joel possesses all (or most) of his responsibility-relevant capacities. Joel is also free in behaving as he does (i.e., there are no counterfactual interveners or devious brain devices waiting to override his decisions). Finally, Joel's awareness falls below an applicable cognitive standard, namely the level of vigilance specified by the norms of vigilance. Not only does the vigilance account offer explanations that conform to conditions on direct responsibility; the vigilance account explains *why* the unwitting agents fulfill these conditions at the time of wrongdoing.

This leaves an interesting question. Why does the vigilance account appear to offer tracing explanations of responsibility? What accounts for the tracing 'residue'? I think two factors contribute to this illusion. First, the normative dimension of vigilance explains why we find people blameworthy (and why we accept judgments of

³² Clarke (2014: 167).

³³ In more recent work, Clarke (forthcoming: 13) suggests that (ii) and (iii) might just be ways of analyzing out the relevant capacities mentioned in (i). The argument here does not depend on whether (i) refers to capacities distinct from the capacities to which (ii) and (iii) refer.

blameworthiness) for unwitting omissions at all. The objection here might conflate this general justification for one feature of our ordinary practice (the norms of vigilance) with the justification for assigning responsibility for unwitting omissions in individual cases (conditions (i)-(iii) mentioned above). The norms of vigilance do not pick out failures that anchor a trace, but justify our general practice of holding people responsible for unwitting omissions.

Second, there might be a conflation of the roles that the cognitive standard plays in the vigilance account. The conditions under which the norms of vigilance apply to an agent are distinct (in this account) from the conditions on blameworthiness that an agent fulfills in virtue of the norms of vigilance. Recall that the application conditions on the norms of vigilance track the possession (or lack thereof) of the vigilance capacity, while the content of the norms tracks the level of attunement that an agent possesses in virtue of exercising vigilance. This adds the illusion of tracing because on any occasion where we assign blame for an unwitting omission we seemingly require an appeal to some earlier time at which the agent fulfills the application conditions of the norms of vigilance and the cognitive standard. Thus, the tracing illusion arises from the confusion of the conditions under which the cognitive standard applies to a particular agent in a context and the conditions under which some agent is culpable for an unwitting omission in those same circumstances.

Conclusion

I began with a puzzle that arises when one adopts a Capacity Account and the theory of moral responsibility that the Capacity Account implies. The puzzle focuses on the fact

that we sometimes hold people responsible even when they appear to lack or fail to exercise relevant agential capacities that are constitutive of moral agency. This appears to indicate that moral responsibility does not track the possession and exercise of capacities. I argued that one should develop the Capacity Account in certain ways to resolve the puzzle. The development occurs along two lines: (i) a distinct normative account of moral responsibility for unwitting omissions that claims that agents are culpable for their unwitting omissions when their awareness is substandard with respect to some applicable cognitive standard, and; (ii) a framework that specifies a set of adequate neural realizers for the abilities posited by the normative account. This framework specifies what the mechanism of control in unwitting omissions comes to in a way that is compatible with the Capacity Account and the corresponding approach to moral responsibility mentioned at the outset.

If this is right, then the central question for picture of responsibility that falls out of the Capacity Account is not whether we have the capacities required for us to be the appropriate targets of responsibility ascriptions, but instead, the conditions under which we fail to meet appropriate standards for vigilance.

References

Amaya, S. 2013. "Slips," in *Noûs* 47:3, 559-576.

Amaya, S. and Doris, J.M. 2014. "No Excuses: Performance Mistakes in Morality," in *The Handbook of Neuroethics*, ed. J. Clausen and N. Levy (Dordrecht: Springer), 253-272.

Aron, A.R. *et al.* 2007. "Triangulating a cognitive control network using diffusion-weighted magnetic resonance imaging (MRI) and functional MRI," in *Journal of Neuroscience* 27, 3743-3752.

Aron, A.R., and Poldrack, R.A. 2006. "Cortical and subcortical contributions to Stop signal response inhibition: role of the subthalamic nucleus," in *Journal of Neuroscience* 26, 2424-2433.

Bogacz, R. *et al.* 2010. "The neural basis of the speed-accuracy tradeoff," in *Trends in Neuroscience* 33, 10-16.

Botvinick, M.M. 2007. "Conflict monitoring and decision making: reconciling two perspectives on anterior cingulate function," in *Cognitive, Affective, and Behavioral Neuroscience* 7, 356-366.

Botvinick, M.M. and Cohen, J.D. 2014. “The Computational and Neural Basis of Cognitive Control: Charted Territory and New Frontiers,” in *Cognitive Science* 38:6, 1249-85.

Cavanagh, J.F. *et al.* 2011. “Subthalamic nucleus stimulation reverses mediofrontal influence over decision threshold,” in *Nature Neuroscience* 14, 1462-1467.

Clarke, R. 2014. *Omissions: Responsibility, Agency, and Metaphysics* (Oxford: Oxford University Press).

Clarke, R. forthcoming. “Ignorance, Revision, and Commonsense,” in *Responsibility: The Epistemic Condition*, ed. P. Robichaud and J.W. Wieland (Oxford: Oxford University Press).

Craver, C. 2007. *Explaining the Brain: The Mosaic Unity of Neuroscience* (Oxford: Oxford University Press).

Doris, J. 2015. *Talking to Ourselves: Reflection, Ignorance, and Agency* (Oxford: Oxford University Press).

Doris, J. and Stich, S. 2007. “As a Matter of Fact: Empirical Perspectives on Ethics,” in *The Oxford Handbook of Contemporary Philosophy*, ed. F. Jackson and M. Smith (Oxford: Oxford University Press), 114-52.

Ekstrom, L. 2000. *Free Will: A Contemporary Study* (Boulder, CO: Westview Press).

Fischer, J.M. and Ravizza, M. 1998. *Responsibility and Control: A Theory of Moral Responsibility* (Cambridge: Cambridge University Press).

Fischer, J.M. and Tognazzini, N. 2009. "The Truth about Tracing," in *Noûs* 43:3, 531-56.

Forstmann, B.U. *et al.* 2008. "Striatum and pre-SMA facilitate decision-making under time pressure," in *Proceedings of the National Academy of Sciences USA* 105, 17538-42.

Forstmann, B.U. *et al.* 2010. "Cortico-striatal connections predict control over speed and accuracy in perceptual decision making," in *Proceedings of the National Academy of the Sciences USA* 107, 15916-15920.

Frankfurt, H. 1998. *Necessity, Volition, Love* (Cambridge: Cambridge University Press).

Ginet, C. 2000. "The Epistemic Requirements for Moral Responsibility," in *Philosophical Perspectives* 14, 267-77.

Graham, P. 2014. "A Sketch of a Theory of Moral Blameworthiness," in *Philosophy and Phenomenological Research* 88:2, 388-409.

Ivanoff, J. *et al.* 2008. “fMRI evidence for a dual process account of the speed-accuracy tradeoff in decision-making,” in *PLoS ONE* 3, e2635.

Jahafri, S. 2011. “Effective connectivity reveals important roles for both the hyperdirect (fronto-subthalamic) and the indirect (fronto-striatal-pallidal) fronto-basal ganglia pathways during response inhibition,” in *Journal of Neuroscience* 31, 6891-6899.

Kerns, J.G. 2006. “Anterior cingulate and prefrontal cortex activity in an FMRI study of trial-to-trial adjustments of the Simon task,” in *Neuroimage* 33, 399-405.

King, J.A. *et al.* 2010. “Post-error behavioral adjustments are facilitated by activation and suppression of task-relevant and task-irrelevant information processing,” in *Journal of Neuroscience* 30, 12759-12769.

Kool, W. *et al.* 2010. “Decision making and the avoidance of cognitive demand,” in *Journal of Experimental Psychology: General* 139, 665-682.

Mulder, M.J. *et al.* 2012. “Bias in the brain: a diffusion model analysis of prior probability and potential payoff,” in *Journal of Neuroscience* 32, 2335-2343.

Nelkin, D.K. 2008. “Responsibility and rational abilities: Defending an asymmetrical view,” in *Pacific Philosophical Quarterly* 89:4, 497-515.

- Nichols, S. 2015. *Bound: Essays on Free Will and Responsibility* (Oxford: Oxford University Press).
- Pereboom, D. forthcoming. "Omissions and Different Senses of Responsibility," in *The Ethics and Law of Omissions*, ed. D.K. Nelkin and S.C. Rickless (Oxford: Oxford University Press), 179-91.
- Raz, J. 2011. *From Normativity to Responsibility* (Oxford: Oxford University Press).
- Rosen, G. 2003. "Culpability and Ignorance," in *Proceedings of the Aristotelian Society* 103, 61-84.
- Rosen, G. 2004. "Skepticism about Moral Responsibility," in *Philosophical Perspectives* 18, 295-313.
- Scanlon, T.M. 1998. *What We Owe to Each Other* (Cambridge, MA: Belknap Press).
- Scanlon, T.M. *Moral Dimensions: Permissibility, Meaning, and Blame* (Cambridge, MA: Belknap Press).
- Shenhav, A., Botvinick, M.M., and Cohen, J.D. 2013. "The Expected Value of Control: An Integrative Theory of Anterior Cingulate Cortex Function," in *Neuron* 79, 217-240.

Sher, G. 2009. *Who Knew: Responsibility without Awareness* (Oxford: Oxford University Press).

Shoemaker, D. 2015. *Responsibility from the Margins* (Oxford: Oxford University Press).

Smith, A. 2005. "Responsibility for Attitudes: Activity and Passivity in Mental Life," in *Ethics* 115:2, 236-271.

Smith, A. 2008. "Control, Responsibility, and Moral Assessment," in *Philosophical Studies* 138, 367-92.

Smith, A. 2015. "Attitudes, Tracing, and Control," in *Journal of Applied Philosophy* 32:2, 115-132.

Strawson, P.F. 1962. "Freedom and Resentment," in *Proceedings of the British Academy* 48, 187-211.

Timpe, K. 2011. "Tracing and the Epistemic Condition on Moral Responsibility," in *The Modern Schoolman* 88:1-2, 5-28.

van Inwagen, P. 1983. *An Essay on Free Will* (Oxford: Clarendon Press).

van Inwagen, P. 1989. "When is the Will Free?" in *Philosophical Perspectives* 3, 399-422.

van Maanen, L. *et al.* 2011. "Neural correlates of trial-to-trial fluctuations in response caution," in *Journal of Neuroscience* 31, 17488-95.

Vargas, M. 2013. *Building Better Beings: A Theory of Moral Responsibility* (Oxford: Oxford University Press).

Vargas, M. manuscript. "Moral Negligence and Social Self-Governance."

Wallace, R.J. 1994. *Responsibility and the Moral Sentiments* (Cambridge, MA: Harvard University Press).

Wiecki, T.V. and M.J. Frank. 2013. "A computational model of inhibitory control in frontal cortex and basal ganglia," in *Psychological Review* 120, 329-355.

Wolf, S. 1990. *Freedom within Reason* (Oxford: Oxford University Press).

Zimmerman, M. 2008. *Living with Uncertainty* (Cambridge: Cambridge University Press).