

VIGILANCE AND CONTROL*

Samuel Murray, University of Notre Dame
Manuel Vargas, University of California, San-Diego

(*Forthcoming in Philosophical Studies*)

Abstract

We sometimes fail unwittingly to do things that we ought to do. And we are, from time to time, culpable for these unwitting omissions. We provide an outline of a theory of responsibility for unwitting omissions. We emphasize two distinctive ideas: (i) many unwitting omissions can be understood as failures of appropriate vigilance, and; (ii) the sort of self-control implicated in these failures of appropriate vigilance is valuable. We argue that the norms that govern vigilance and the value of self-control explain culpability for unwitting omissions.

Introduction

Consider the following case:

Bourbon. As Randy is about to leave his home for the office, Al calls to tell him that they are out of bourbon. His regular route to the office takes him right by a liquor store, and Randy tells Al he'll buy some. Between his home and the liquor store, Randy starts thinking about a paper he is writing on omissions. He continues thinking about his work until he arrives at the department, where he realizes that he has forgotten the bourbon.¹

Randy's failure to bring the bourbon is criticizable. When Al can't make the Old Fashioned he's been craving all day, he'll blame Randy. The structure of this case generalizes to other kinds of omissions. Caretakers forget their kids in the backseat of cars (see Amaya, 2013) or fail to observe that they've left the stove on (see Ahrens, 2016). Less

* The authors take equal credit for the ideas in this paper, though the second author notes that the first author deserves most of the credit for writing. Versions of this paper were presented at the 2015 Stanford Graduate Student Workshop, the 2015 Philosophy and Science of Self-Control Conference, and the Moral Attitudes and Decisions Lab at Duke University. We thank the audiences for helpful comments, especially those at the Philosophy and Science of Self-Control Conference. Special thanks to Al Mele, Adina Roskies, Katrina Sifferd, and Chandra Sripada for many conversations on different parts of this project. Finally, special thanks to Walter Sinnott-Armstrong and Santiago Amaya for providing written comments on multiple drafts of this paper that helped us correct many deficiencies. We also want to acknowledge the insightful criticisms of an anonymous referee at *Philosophical Studies* that helped us make this paper much clearer and sharper. Research for this paper was supported by a Philosophy and Science of Self-Control grant from Florida State University awarded to both Samuel Murray and Manuel Vargas. The views expressed in this paper do not reflect the opinions of the other grantees or Florida State University.

¹ Clarke (2014: 164) is the inspiration for the example. Mele (2006), Murray (2017), Graham (2014), and Sher (2006) discuss similar cases.

dramatically, people forget their keys, forget about meetings, miss exit signs on the highway, and forget anniversary or birthday dates. Ordinary practice suggests that many instances of these failures merit condemnation or fault-finding.

The practice is not puzzling in itself; rather, puzzles arise when one attempts to map practice to theory. The problem is that in *Bourbon*—and relevantly similar cases—the offending agent’s conduct lacks all familiar actional and valuative antecedents that might ground responsibility. There is no decision, volition, intention, belief, desire, choice, or judgment (among other things) that seems to explain the omission.

In what follows, we offer part of an outline of a general theory of responsibility for these failures.² We’ll focus on *Bourbon*, but we think that our account generalizes to other cases of unwitting omission. Additionally, we explain why people are motivated to take responsibility for some of their unwitting omissions.

Our account emphasizes several distinctive ideas. First, many unwitting omissions can be understood as failures of appropriate vigilance, i.e., failures to meet some standard of information management, given an ability to meet these demands. However, this vigilance-focused account raises a puzzle of its own: why are we typically concerned to meet sometimes demanding norms of information management? A second idea comes in here, namely that self-control of the sort implicated in meeting these standards is of particular value to us, both individually and socially. Together, these features account for

² The notion of responsibility is a multi-headed beast. In this paper, we’re interested in the accountability sense of responsibility, where an individual is responsible in the accountability sense only when the individual’s conduct makes it appropriate to resent or be indignant toward that individual (see Shoemaker, 2015). Hence, we are interested in a kind of responsibility implicated in making it appropriate for individuals to bear certain Strawsonian reactive attitudes toward one another. In what follows, our use of the term ‘responsibility’ tracks this phenomenon. While there are interesting issues surrounding alternative kinds of responsibility (such as attributability, answerability, legal responsibility, etc.), we set these aside to focus on accountability. Thus, the use of ‘responsibility’ functions as shorthand for ‘the accountability kind of responsibility’.

the basic structure of explaining *Bourbon*-style cases, and it provides a framework readily amenable to further elaboration in terms of, for example, specifically moral or other forms of blame and fault-finding.

In Section 1, we briefly define unwitting omissions and discuss some issues with standard explanatory strategies applied to responsibility for these omissions. In Section 2, we sketch the details of recent empirical work on the nature of cognitive control. This empirical work gives us important insights into the capacity for information management. In Section 3, we connect these empirical details to an account of vigilance and specify the role that failures of vigilance play in *Bourbon*-style cases. In Section 4, we explain the distinction between culpable and non-culpable unwitting omissions that falls out of the vigilance account. In Section 5, we explain what determines the standards of appropriate vigilance and why agents agree to live up to these demanding standards. We also claim that our account of responsibility in *Bourbon*-style cases generalizes to all cases of unwitting omission.

1. *Unwitting omissions*

The term ‘unwitting omission’ refers to failures to act where the agent is unaware of the fact that she will fail and is failing to act during the duration of her failure. Unwitting omissions are a species of unintentional omissions, where an unintentional omission is any failure of an agent to act or refrain that is not the upshot or result of a relevant intention. An unwitting omission is any unintentional omission where the agent is unaware at every relevant moment that she will omit and is omitting.

Ordinary practice suggests that people are sometimes responsible for their unwitting omissions (cf. Murray et al., Forthcoming). This practice does not sit well with

control-based accounts of the grounds of responsibility, which are perhaps the most common approach to moral responsibility. On control views, attributions of responsibility are justified in virtue of an agent's control being implicated in bringing about the state of affairs for which she is responsible.³ The problem is that people do not appear to control their unwitting omissions. Omissions of this sort do not stem from decisions or choices; instead, these omissions seem to stem from the *absence* of decision or choice.

Responsibility theorists have responded to this problem in several ways. Some attempt to explain responsibility for unwitting omissions by tracing the occurrence of some unwitting omissions back to a controlled action that then grounds responsibility for the subsequent unwitting omission.⁴ Despite several sophisticated proposals, many suspect that tracing accounts of responsibility are inadequate. No proposal seems to offer a fully general account of responsibility (Agule, 2016; Shabo, 2015), and the view seems to misconstrue the phenomenology of responsibility attribution (Graham, 2014). Other theorists, operating within the control-based tradition, argue that responsibility-apart unwitting omissions violate some normative epistemic requirement. The offender should have been aware either of her omitting or her risk of omitting. This strategy has been criticized for failing to explain how an epistemic violation grounds distinctly moral appraisal (Sher, 2009: 81-82). Additionally, the view comes with the cost of gesturing at some unspecified capacity governed by the normative epistemic requirement.

The combined conceptual problems and metaphysical baggage have pushed some to abandon the control-based framework entirely. Some explain responsibility for

³ This formulation presumes that the objects of responsibility are states of affairs. However, nothing substantive hangs on this presumption. We could easily translate this into an action-centered view, where an agent is responsible in virtue of her control in acting so as to bring about a certain state of affairs.

⁴ For applications of tracing strategies to the problem of responsibility for unwitting omissions, see Timpe (2011) and Fischer & Tognazzini (2009).

unwitting omissions in terms of the agent's manifesting an objectionable lack of care (Björnsson, 2017; Sripada, 2016; Mason, 2015). However, the Quality of Will views have difficulty explaining responsibility for unwitting omissions. Some argue that individuals cannot be responsible for unwitting omissions precisely because these individuals do not manifest a deficit of care in omitting unwittingly (cf. Sripada, 2015; Rosen, 2008). All of these problems have motivated skepticism about the possibility for responsibility for unwitting omissions. This option becomes more attractive in light of the aforementioned difficulties. One might accept these revisionary conclusions (e.g., Doris, 2018), but we think there is a more appealing route forward.

In what follows, we present a modified control-based theory that explains responsibility for unwitting omissions. Consider this an acknowledgement of partisan flag-planting. Our hope is that in generating an independently motivated reconstruction of the control condition on responsibility, we can devise a solution to the problem of responsibility for unwitting omissions that avoids metaphysical hand-waving and deflates some of the motivation for adopting alternatives to the control-based theory of responsibility.

2. Cognitive control and information management

As mentioned in Section 1, explanations of responsibility in *Bourbon*-style cases generally begin with some account either of a modification to the awareness condition on responsibility or of some way in which Randy (or agents in similar circumstances) displays objectionable levels of care in omitting (cf. Sher, 2009; Clarke, 2014: 167; Harman, 2011; Guerrero, 2007: 80; Björnsson, 2017; Smith, 2005).

We want to begin in a different spot. On our approach, we start with what we know about our cognitive architecture, and in particular, the challenges of information management in human beings. From there, we build up to a conceptually robust story about responsibility-relevant forms of control. Then, we discuss the kinds of demands we can reasonably place on creatures with cognitive architecture like ours. On our account, Randy is culpable for his unwitting omission in virtue of failing to be appropriately vigilant. Failures of appropriate vigilance are failures to live up to a reasonable demand for information management, given the agent's capacity to do so. In this and the next section, we specify what the relevant capacity is and (in Section 5) what these demands for information management amount to.

Human beings are finite creatures with limited cognitive capacities. Nevertheless, human beings must navigate information-rich environments with these strict capacity limitations.⁵ There is no way to apprehend every bit of available information, so the brain must determine how to allocate cognitive resources to processing certain kinds of information and selecting certain actions at the expense of others. This allocation issue presents the brain with two problems given information overload. One is a problem of synchronic information processing. Given a variety of immediate conflicting demands, how does the brain pick out salient or relevant bits of information? The other problem is one of diachronic information processing. Given a variety of plans and commitments, how do we structure our activities in a way that allows us to act efficiently without straining our severely limited memory or attentional capacities?

⁵ Despite controversy about the *source* of these limitations, nobody disagrees that there are such limitations. Try as hard as you want, you'll never be able to do two two-digit math problems simultaneously. Whether this is a matter of limited metabolic resources, functional limitations, structural limitations, or whatever else is irrelevant to our account.

One way to work around these limitations is to use heuristics that efficiently process information in a form that is both ecologically bounded and typically effective. There is now strong evidence that our cognitive processing includes a ‘fast’ or ‘hot’ system that rapidly processes information and utilizes certain ‘automatic’ or prepotent responses to certain informational inputs (Evans and Stanovich, 2013). However, there will be times when the prepotent response is suboptimal. In these instances, a ‘slow’ or ‘cold’ system is needed to override the automatic system and control the flow of information across alternative processing pathways toward some other response. Thus, in situations where the automatic response is less than optimal (or perceived to be less than optimal), the cold system can induce regulation over information processing that allows for flexible control (cf. Cohen, 2017: 7-10 for an overview).

There is good evidence that the brain’s computational labor divides along these lines, supporting two different processing modes. One mode is habitual and supports automatic responses, while the other accounts for slower, controlled responses. This division is an optimized arrangement for processors that need to balance speed and accuracy efficiently (Kool & Botvinick, 2014). Automatic processing provides speed, though the presence of novel stimuli or unique situational demands points to the need for the self-control that controlled processing provides.⁶

⁶ Many philosophers likely understand this division of computational labor as the division between System-1 and System-2 processing (sometimes labeled the Dual Process Theory). While the label is useful, many cognitive scientists question the System-1/System-2 framework and the presuppositions of that framework. Since the original proposal of the two-system division over 40 years ago (Posner and Snyder, 1975), the hypothesis has been criticized and significantly revised (see Cohen et al., 1990 for overviews; see Cohen, 2017 for criticisms). These advancements and revisions, however, do not affect the points that we make below. Our view does not depend on any substantive formulation of the Dual Process Theory or a specific view on the mechanisms or circuits that realize these information-processing structures *or* the sorts of processes that count as either habitual or controlled are irrelevant to our argument.

The cognitive control system is the main hub of controlled processing (Cohen, 2017: 4). Roughly, the cognitive control system monitors conflicting demands that are made on various cognitive resources. In addition to monitoring, the cognitive control system also induces regulation on lower-level processing systems by biasing certain processing pathways over others.

To see how this works, consider a variant of the *Bourbon* example, one where Randy remembers to get the bourbon. Randy's various motor movements result from automatic responses given the goal of getting to the office. When he sees the liquor store, there are conflicting demands: the automatic response to continue moving toward the office and the cognitively mediated desire to buy bourbon. Because Randy (in the "success" variant of the *Bourbon* case) wants to get the bourbon for Al, the cognitive control system needs to kick in to bias the processing pathways that favor the non-automatic response.⁷

The issue of managing conflicting demands shows why it is useful for organisms like us to have a controlled information-processing system, but it also highlights the important functions of the cognitive control system. First, the control system monitors conflicting demands from lower-level processing units. Second, the control system stores task representations and implements the appropriate one given both expectations about future reward balanced against current costs of implementation *and* relevant environmental and situational cues (Shenhav, Botvinick, and Cohen, 2013). This last component of cognitive control is important for cases like *Bourbon*.

⁷ Sometimes, as they say, it's better to be lucky than to be good. But we want to set aside cases where Randy gets lucky and has the relevant thought in some way that is disconnected from any interesting notion of cognitive control. We suspect there are interesting things to say if he has the thought because he is hit by a stray cosmic ray, or because of the workings of a bourbonic thought induction field of a mad scientist, or even the more pedestrian case of a mechanism disconnected from cognitive control. Whatever the proper verdict in those cases, however, they are not the phenomenon that is of interest to us.

For example, consider the successful variant of the *Bourbon* case. While walking, Randy is thinking about his paper. When he approaches the store, he needs to shift his attention away from the paper (at least for some time) and toward the store. Thus, the cognitive resources needed to focus on the paper conflict with the resources needed to notice the store and get the bourbon. The cognitive control system monitors this conflict and signals for a new task representation to be activated. The appropriate task representation is activated and Randy walks into the store. Hence, the cognitive control system is needed for activities just like this one: some new stimulus creates a conflict that demands the activation and maintenance of a new task representation.

This means that cognitive control is a crucial constituent of agential control (see Murray, 2017: 517-19). Our sketch of the functional profile of cognitive control glosses over a mountain of mechanistic and computational details that have come out over the last decade or so. The monitoring, specification, and implementation functions can be described in precise computational terms. Many of these computations are derived from reinforcement learning paradigms that explain the various components of cost/benefit decision-making (Kool, Gershman, and Cushman, 2017). Recent neuroimaging work has mapped these computational procedures to underlying mechanisms, such as the dorsal anterior cingulate cortex, dorsolateral prefrontal cortex, basal ganglia, brain stem dopaminergic nuclei, subthalamic nucleus, and locus coeruleus (see Brown, 2017 for an overview of the neuroimaging work). For our purposes, however, these details—while interesting—are not critical to our discussion of the control condition on responsibility.

The previous example simplifies the way in which the cognitive control system balances conflicting demands on our limited cognitive resources. In addition to monitoring for these sorts of demands, the brain balances various future-oriented

commitments that we cannot fulfill all at once. The cognitive control system is also engaged in balancing prospective memory with current actional and environmental demands (Einstein & McDaniel, 2005). Randy has to work all day, get to the liquor store, and find his way to the office (among, plausibly, a number of other things). He cannot do these all at once, so he plans to do some things later. These prospective memories are encoded by the cognitive control system and bound to certain cues the detection of which causes signals to be sent that bring the relevant task representations online.

Thus, the cognitive control system solves the problem of synchronic and diachronic information management mentioned at the outset. These structures and their functions point to our unique abilities to balance a variety of competing demands in the face of strict limits on attentional and memory capacity. Understanding cognitive control at the functional level also tells us something important about the way in which we manage informational demands and structure our plans in light of cognitive limitations. Given this sort of neural architecture, we claim, there is good evidence that Randy has a robust capacity for information management that is implicated in his being called to account for his unwitting omission.

Some may worry that our account of cognitive control fails to connect up with accounts of agential control in philosophically interesting ways. In order to successfully recall his plans, for example, Randy's cognitive control system must kick in. But that seems to invite the question of what role *Randy* himself plays in controlling his cognitive control. However, we are offering a partial componential analysis of agential control. The cognitive control system partially constitutes agential control. For this kind of deflationary account of agency, then, it does not make sense to ask whether Randy controls his cognitive control. Our commitment to this deflationary picture stems from assuming a

causal theory of action that is compatible with naturalistic accounts of human agency. This, we think, is a package of commitments that most control-based theories can accommodate (and perhaps willingly embrace). While some may disagree with these views of action and agency, those disagreements occur at a level outside the scope of this paper.

3. *Vigilance*

Our claim is that Randy is culpable for his unwitting omission in *Bourbon* because he fails to live up to certain reasonable demands for information management *and* he has the capacity to meet these demands. Thus far, we have sketched the lower-level cognitive architecture that, we think, underpins the sort of capacity that is at issue in our fault-finding in cases of omission.⁸ We take it that the relevant capacity is vigilance, and that Randy's failure of vigilance is partly what explains his responsibility in *Bourbon*. In this section, we explain what the role of vigilance is in our account.

Vigilance is a psychological capacity that is partially constitutive of the socially salient ability to manage informational demands. Vigilance is a capacity to become occurrently aware of morally or prudentially relevant considerations for acting. This capacity, then, is crucial for navigating the sorts of information-rich environments that moral agents find themselves in.

⁸ In what follows, we build our account on the presumption of a species-typical neural system. There could be either modest departures from this species-typical set-up or localized functional deficits at the neurobiological or mechanistic level that raise interesting questions about how culpability interacts with shifting environmental conditions. A full theory should have something to say concerning agents with non-typical cognitive control mechanisms or alternative computational bases for vigilance. However, our interest is in ordinary, paradigmatic cases of unwitting omissions, so we are (reluctantly) putting aside those intriguing varieties of cases.

The details about the cognitive control system tell us something important about the functional profile of vigilance. Recall that the cognitive control system accounts for awareness of relevant environmental factors and sensitivity to those environmental factors. We take it that vigilance has a similar functional profile, namely that vigilance enables agents to be aware of and sensitive to relevant considerations in a context. Thus, the account of cognitive control informs the account of vigilance.⁹

Vigilance plays two important roles in an agent's self-governance. First, it enables agents to be aware in the sense of directing attention to action-relevant aspects of one's environment. Second, it enables an agent to be sensitive (to a greater or lesser degree) to those relevant aspects in a way that is crucial for practical deliberation.

This psychological capacity is structured by an agent's valuative commitments, knowledge, experience, and other dispositions. Having vigilance enables the agent to form certain action-relevant representational or motivational states in light of attending to relevant considerations (or, in some cases, vigilance enables the agent to manifest certain dispositions to form such action-relevant states in appropriate circumstances).

Vigilance enables Randy to regulate information in ways that come up to standards of information management that govern him in light of his various commitments. And possessing vigilance is part of what makes him culpable for not getting the bourbon. Thus, because Randy has vigilance, he can be held to account for his failures to meet certain demands of information management (supposing, for now, that these demands are reasonable).

⁹ There could be alternative, bottom-up kinds of processing relevant to vigilance (see Gadziola & Wesson 2016). What is important for our account is that there is at least one plausible, coherent story about the realization of vigilance in some kind of neural machinery. We do not suggest that it is the only story or the full story. So, for example, this is compatible with the idea that agents that lack species-typical neural architecture—or have completely different sorts of neurobiological structures—can be vigilant.

There are two components of appropriate vigilance. First, the agent manifests a degree of vigilance whereby the agent will detect and properly manage various informational demands within a given context. Second, the agent's manifestation of vigilance satisfies some reasonable demand for a certain level of vigilance. Conversely, a failure of appropriate vigilance is one where the agent has vigilance, but fails to manifest vigilance in a way that meets certain demands. This two-part account explains a bit of ambiguity in ordinary vigilance talk. Sometimes, when we talk about being vigilant, we are simply referring to the level of vigilance that the agent manifests. In these cases, 'vigilance' functions more like a descriptive term where agents can be vigilant to a degree but fail to detect relevant, accessible considerations. On the other hand, being vigilant can also refer to the agent's manifesting a degree of vigilance that meets a reasonable demand for vigilance. In these cases, 'vigilance' functions more like a success term.

The fact that vigilance comes in degrees introduces the possibility of semi-vigilance. For instance, suppose that, in a variant of *Bourbon*, Randy promises to bring bourbon to the office party. However, when he gets to the store, he remembers to get a brown alcohol, but he thinks of scotch rather than bourbon, and so he buys the wrong thing. Is Randy criticizable in this case? The issue is complicated, but our approach suggests some resources for thinking about the case.

Many of the tasks we perform are elaborate and composed of various sub-tasks. Further, task demands can be more or less fine-grained. This implies that there are varying levels of breakdown in task performance associated with sub-task performance and task performance at the appropriate level. Demands on task performance vary contextually, affecting the perceived severity of these breakdowns. If the stakes dictate that Randy purchase a brown spirit, then scotch might do—even if it isn't Al's favorite.

But if the circumstances demand Blanton's, Randy's bringing scotch might be a more serious transgression. Context shapes responsibility for cases of semi-vigilance.

Some instances of responsibility-apt vigilance might seem susceptible to problems of causal deviance. For instance, consider a different variant on *Bourbon*, one where Randy (falsely) believes that he needs to get scotch for the party, but through inattentiveness, buys bourbon instead. Is this a failure of appropriate vigilance? Much of the story about responsibility here will turn on how Randy came by his false belief. But supposing there's no failure of appropriate vigilance to account for his false belief, then we think Randy is not criticizable in this case of deviance. This, however, has less to do with Randy's vigilance and more to do with the fact that there's no occasion for criticism in the case where Randy shows up at the office with bourbon.

We claim that part of what makes Randy responsible for his unwitting omission is a failure of appropriate vigilance. That is, there is a failure of information management given a demand to properly manage information. The discussion of vigilance and cognitive control gives us good reason to believe that Randy has the sort of capacity needed to meet these demands. In the next section, we argue that a plausible story of the difference between culpable and non-culpable unwitting omission falls out of the vigilance account.

4. *Culpable and non-culpable unwitting omission*

To see the difference between culpable and non-culpable omission, compare the *Bourbon* case to a different omission case. Suppose that

Lottery. Faye is down on her luck, broke, and recently unemployed. Faye needs money to support her family. She passes by a store while walking down the street. Had she walked into this store, she would have been the

millionth customer and won a huge cash prize.¹⁰ Unfortunately, from Faye's vantage point, she cannot discern this. It seems that Faye is not criticizable for failing to win the prize (even though she could easily do so) because her successfully doing so would be a matter of mere luck.¹¹ And what makes it the case that her winning is a matter of mere luck is that she is ignorant about the contest.

The difference between Randy and Faye is that Randy, but not Faye, *should have been more vigilant*. This implies that others can justifiably criticize Randy for substandard information management. This does not seem true of Faye. Perhaps for some sense of 'should' it is the case that she should have been more vigilant. But that use of 'should have been vigilant' seems importantly different from the use that describes Randy's situation.

Perhaps the most plausible story about the norms at issue (part of what makes it true that Randy *should* have been more vigilant) is to understand them as special obligations that derive from commitment-oriented relationships. Both Randy and Faye have special obligations to their friends or family in virtue of being in certain relationships with those people. Those relationships entail adopting certain commitments oriented toward the other person in the relationship. Some of these commitments entail obligations to direct attention or privilege the distribution of attentional resources toward caring for the other person in the relationship. Thus, Randy, in virtue of being a friend, has certain commitments oriented toward Al that include being obligated to attend to Al in particular ways. Faye has these commitments as well, but certain properties of the

¹⁰ This is based on a case in Zimmerman (1986: 205).

¹¹ The inference here relies on a simple blame-implies-can principle (cf. Buckwalter & Turri, 2015; Henne et al., 2016). If S is blameworthy for A-ing, then S can avoid bringing it about the A-ing occurs in virtue of her agency. If S can avoid bringing it about that A occurs in virtue of her agency, then whether A-ing occurs in virtue of her agency is not a matter of luck. Hence, if A-ing occurs as a matter of luck, then S is not blameworthy for A-ing. We rely on the idea that individuals cannot be held to account for occurrences outside their control. If an occurrence is a matter of luck, then whether there is such an occurrence lies beyond the agent's control.

epistemic context releases her from these obligations. Randy's epistemic context does not share these properties, and so he remains bound to these special obligations to Al. This is what makes it the case that Randy *should* have been more vigilant.

Faye needs to get money to support her family, but she cannot determine that walking into the store would help her achieve this goal. The upshot of the case is that Faye is not obligated to walk into the store, but this is because of a defect in the epistemic context. Thus, it is not the case that Faye *should* have been vigilant because she *could not* have been vigilant.

We take it that Randy *could* have known better (in some relevant sense of 'could') because he has vigilance. By this, we mean that Randy is able to generate or come to have action-appropriate attitudes. Randy has this ability in virtue of possessing vigilance, and that capacity is subserved by the usual sorts of cognitive control structures.

From the standpoint of ordinary talk about capacities and abilities, this should seem unremarkable. Alas, accounts of capacities are as fraught as anything in philosophy. For our purposes, we'll help ourselves to two assumptions: (1) there are useful ways of talking about capacities that are compatible with most ways of understanding the ontology of action, and; (2) our talk of capacities or abilities is disjointed, reflecting different practical and theoretical interests. Some notions of capacity reflect an interest in explanation and prediction. Other notions of capacity reflect different concerns—e.g., standards of evaluation in a social practice. Whether a dam is capable of holding back a flood, for example, involves a notion of capacity structured by an interest in prediction. Whether someone has the ability to meet a deadline for submitting a paper involves more than prediction. Indeed, one might predict that this particular author is exceedingly unlikely to turn the paper in on time. In each case, the relevant notion of capacity turns

on facts about the target of evaluation (the dam and the damned author). However, each involves a different set of concerns that pick out some different notions of capacity. Our presumption is that where there are interests in coordinating and organizing social practices, talk of an agent's capacities will oftentimes come apart from the kinds of capacities we will attribute to phenomena where our only interest is in prediction and explanation. We trust that both assumptions—that talk of capacities is compatible various ways of carving up the ontology of action and that capacity talk is disjointed in light of diverse practical and theoretical interests—are plausible enough on their face. Anyway, recapitulating familiar debates in the philosophy of action and some versions of the free will problem would take us far afield, so we set those concerns aside. With respect to the disunification thesis, what follows may resolve some questions.

Recall that manifestations of vigilance consist in, for example, the agent's coming to have relevant attitudes formed in light of the agent's attending to relevant considerations. This helps to explain part of the story about Randy. Randy fails to manifest appropriate vigilance. He could have formed (but failed to form) some belief that he needed to get liquor or a desire to go into the store in light of the consideration that he promised to get some for Al. This consideration is the object of Randy's vigilance because it is relevant for him in this circumstance.

In contrast, Faye cannot manifest appropriate vigilance because the environment is not configured in a way that supports vigilance. Randy, on the other hand, is situated in an environment that supports vigilance, because the normatively relevant features of the environment are accessible to Randy. That is, because Randy has vigilance, he can deliberately or spontaneously bring to mind the normative reasons that rationalize his acting in accordance with his obligations toward Al. This makes clear what it means to

say that Randy's reasons are *accessible* to him. An agent that is ignorant with respect to normatively relevant features of some situation can access the reasons given by those normatively relevant features in virtue of being vigilant with respect to those normatively relevant features of the situation.

We need to be careful not to over-generalize the features of vigilance, however. Someone might wonder why Faye cannot be vigilant with respect to considerations that bear on the cash prize. The reason she cannot, we think, is that there was never a time (we suppose) when she *had* the relevant information or was in a position to gain the relevant information. Faye does not exhibit a failure of information maintenance that seems crucial to vigilance. Randy, on the other hand, does exhibit a failure to maintain relevant information. This maintenance does not require Randy to hold the relevant information within the scope of occurrent awareness; rather, information maintenance requires appropriate memory encoding and utilizing cues that increases the likelihood of recall at the right time (e.g., writing a note, setting a reminder on your phone, etc.).

At this stage, the issue of moral luck appears to rear its head. After all, aren't the subpersonal occurrences necessary for vigilance a matter of luck? And, if they are, how can Randy be responsible for failures of vigilance? We accept that there is an interesting challenge here, and that in some form or another the question can be pressed against nearly any deflationary, causal, control-based account of responsibility.

Like most control theorists, we don't think luck is fatal to responsibility. Control theorists can and have addressed the challenge of luck (e.g., Vargas, 2012; Hartman, 2016). Second, we are inclined to think that the discussion in the next section of the *value of self-control* can be understood as highlighting a good that (in the ordinary case) can only be attained if we tolerate some amount of luck. That said, given that our present ambition

is to tell a story about unwitting omissions that is continuous with control-based theories of responsibility, we are content to let our account ride or fall on the more general resources of control-based accounts for addressing questions of luck.

Here, then, is where we are: the notion of vigilance helps to explain why Randy is culpable for his omission in *Bourbon*. This also explains what happens with other cases of unwitting omission. Randy's case highlights certain dimensions of the monitoring aspects of vigilance, though other kinds of unwitting omission highlight different dimensions. For instance, failures of vigilance might explain underbelief, false belief, or failures of inference. And when these failures of vigilance are connected to the special obligations that derive from commitment-oriented relationships, culpability for these failures will have a similar explanation to the one described above.

5. *Competence and Self-Control*

We have, thus far, explained the way in which Randy possesses a capacity relevant to meeting certain demands. Because Randy has vigilance (but fails to manifest appropriate degrees of vigilance), he can be held to account for his omission. The core feature of this account so far appeals to vigilance, a person-level capacity that enables control over monitoring, implementing intentions, and regulating temporally extended action.

An important puzzle remains: What makes it the case that certain demands for information management are reasonable? And why is it that agents accept that standard as binding on their conduct? When Randy fails to be vigilant, presumably there is a story to be told about his failure, a story that reflects the operations of his neurological, cognitive, and affective architecture. But why view the outcome of those things—things that are, let us suppose, determinative in Randy's failure—as licensing *criticism* of Randy?

A satisfactory account must specify why the demands implicitly encoded in our responsibility ascriptions of unwitting wrongdoers are reasonable. We think that an answer is forthcoming if we focus on the *value of self-control*.

Many philosophers have noted that self-control or self-governance is central to responsible agency (Ekstrom, 2000; Roskies, 2010; Fischer and Ravizza, 1998; Vargas, 2013a). On many such accounts, what makes an individual a morally responsible agent is (at least in part) the general capacity to direct oneself in light of distinctively moral reasons for acting. Perhaps the standard way to flesh out this general capacity is in terms of an agent's (or some mechanism's) responsiveness to reasons (Fischer and Ravizza, 1998; Nelkin, 2011; Roskies, 2012; McKenna, 2013; Vargas, 2013a). According to these accounts, reasons-responsiveness is at least necessary (if not sufficient) for responsible agency. But why is self-control important in the domain of responsible agency?

Put simply, self-control matters because it concerns the ability to conform one's behavior to norms that arise as a result of one's conception of the person one is and hopes to be (and in turn that others reasonably expect one to be). In cases where one is who one hopes to be, two things are true: one has certain competencies that (at least partly) constitute one's being that person and one reliably manifests those competencies in one's behavior. So, if one is (and wants to be) a loving parent, one has the ability to be sensitive to the needs of one's child and the wellbeing of that child and one reliably manifests or acts on those abilities. Self-control matters since it reflects the extent to which one actively seeks to, and succeeds in, conforming behavior to the self that one aims to be, and thus (to some extent) is. In the remainder of this section, we expand this account of the relationship between self-conception, competence, reliability, and self-control.

One's self-conception includes elements such as the relationships that one deems important, the projects that one finds fulfilling, and the activities that one esteems.¹² Despite localized or temporary fluctuations, in the ordinary case, one's self-conception at a time (or over a minimal range of times) expresses many of the values, carings, and aspirations of a given agent at that time (or across that range of times).

Part of the content of one's self-conception is an awareness of special skills or capacities that are necessary to manifest one's overall scheme of cares. These capacities are important in at least three ways (cf. Schapiro, 2001). First, they play an important relational function. Our distinctive capacities relate us to the world in a way that reflects the kind of person we aim to be. Second, they play an instrumental role, in that exercises of these capacities bring about the activities that we esteem and realize the projects that we deem fulfilling. Third, our capacities play a revelatory role, in that exercising our capacities reveals to others certain features of our self-conception. As such, our capacities play a crucial role in helping us live up to and realize our aspirations and ambitions.

Such capacities require or perhaps define a sphere of competence that ranges over a set of actions in which we can reliably engage.¹³ Thus, in virtue of one's self-conception one takes oneself to have certain capacities that one can reliably exercise to bring about certain outcomes. If, for instance, part of one's self-conception includes being a reliable

¹² Of course, one's self-conception will include much more than this. Perhaps one's self-conception divides into normative and descriptive elements. If that is the case, then the aspects of one's self-conception that are important for this paper are the normative elements. Also, the notion of 'self' at issue here likely varies cross-culturally, as Heine, Lehman, Markus and Kitayama (1999) suggest.

¹³ Amaya & Doris (2014: 258-260) also discuss an agent's "zone of secure competence" with respect to similar cases (what we call 'unwitting omissions' they call 'performance mistakes'; as far as we can tell, performance mistakes appear to be a subclass of unwitting omission). However, Amaya and Doris claim that normative competences constitute the zone of secure competence, and equate these normative competences with appropriately reasons-responsive mechanisms. Our discussion of competence provides a fuller characterization (we think) of what these normative competences are, how they link up with other aspects of our psychology, and why these competences are considered normative.

teacher or a reliable parent, then some of the act-types that fall within one's sphere of competence will include activities that fall under those roles.

Competence, here, is a function of the degree of reliability that one manifests in undertaking certain activities. Thus, there is an important relationship between self-conception, competence, and reliability. The central connection is that we want to display reliability in those domains where we value competence (where it fits into our ideal self-conception). Our valuing such competence also signals to others that we desire to be assessed in light of the norms that attach to occupying particular roles, the expectations that these norms generate, and the demands from others that they invite.

When people make mistakes (and unwitting omissions are a subset of mistakes), they violate demands that emerge because the agent occupies certain roles. Where an agent regards the demands as unreasonable, the demands can be protested. Plausibly, demands are unreasonable in at least two circumstances. The first is when the agent defensibly resists occupying that role. If I am terrible at navigating, then my travel companions have no business expecting me to lead them on a tour through a foreign city. If they have such expectations, and blame my failure of leadership, I can reasonably object to their placing me in such a role. The second is when one's circumstances contain systematic, competence-defeating features. If I am in fact a good navigator, but I falsely believe we are in Brussels rather than Amsterdam, or I am being coerced at gunpoint to lead the group down a blind alley, then my competencies are being defeated or undermined by external forces. Other things equal, the unreasonableness of a demand in either sort of circumstance gives me an excuse for certain kinds of failure.¹⁴

¹⁴ Of course, there will be some question as to how we sort out the excusing kinds of ignorance from the non-excusing kinds, though some of that machinery will derive from a more systematic account of the

Where the demands are reasonable, however, there are two primary responses to mistakes that are available. Either the community must adjust its expectations of the agent (e.g., I am not a good navigator after all) or the agent merits blame (Higgins, 1997). There is oftentimes a high social cost to signaling to others that one lacks the relevant competence. So, individuals are ordinarily motivated to live up to these expectations, and to accept negative assessments in light of failures to meet these expectations (Miller & Monin, 2016). Accepting a negative assessment is an important signal to the community. In a costly fashion, it communicates that one takes oneself to have the relevant competence that figures in the demand, even if one did not (or could not) successfully exercise the capacity in the criticized case (cf. Sinnott-Armstrong, 1985).

This helps us to see the personal significance of competence. Because agents desire to exhibit competence, individuals have an interest in manifesting reliability in those domains they find valuable. Put another way, agents have an interest in manifesting their reliability to themselves and others for the sake of demonstrating that they can be considered reliant. When individuals make mistakes in domains of personal importance, there is a *prima facie* threat to that individual's self-conception in virtue of the threat to that individual's claim to reliability in a particular domain. One way to defuse this threat is to take responsibility for mistakes. In taking responsibility (either moral or non-moral) for some mistake, an individual signals both to herself and others that she possesses the relevant competence (Crocker & Wolfe, 2001). Acknowledging a failure to exercise this competence, even when it is seemingly costly, nevertheless affirms the agent's competence

norms of vigilance. In spite of this problem, it remains true that there are relatively clear-cut examples of ignorance that make a demand unreasonable and thereby furnishes the agent with an excuse.

in that domain. Hence, the sphere of competence helps to explain why individuals would agree to take responsibility for their mistakes.

Importantly, mistakes only threaten one's self-conception (and thus become something for which an agent might take responsibility) when the mistake threatens the claim to reliability. If an agent already has a solid reputation for reliability, then individual mistakes will not constitute a threat. For example, if Christen Press makes a bad run or LeBron James throws a bad pass, nobody will question their competence as professional athletes (readers can feel free to substitute in the names of their favorite professional athletes to test these intuitions). Systematic displays of unreliability or relatively grave singular lapses, however, will threaten even the maximally competent agent (e.g., don't miss *all* your writing deadlines or word might get out that you aren't, in fact, competent at participation in this kind of endeavor!) At bottom, then, the agent's claim to reliability is the foundation for the normative account of expectations and demands that we offer here.

An agent's sphere of competence—whether the agent is a master of that domain, as it were—is crucial for self-control. As we see it, one possesses self-control to the degree that one's behavior falls within one's sphere of competence. By extension, one counts as a responsible agent only to the degree that one is acting within one's sphere of competence. This implies (virtuously, we presume) that self-control and responsible agency are highly context-sensitive and shift according to one's movement into and out of one's sphere of competence (Vargas, 2013b).

All of this supports our contention that self-control plays a valuable role in our lives. These considerations lead us to a normatively significant account of self-control in terms of the possession of capacities that figure into our self-conception and constitute a

sphere of competence. On this account, we are in control of ourselves and what we do simply in virtue of possessing these capacities in a context where nothing blocks or significantly impedes the proper exercise of these capacities.¹⁵

In this way, we think that the competence account explains why certain demands for information management are reasonable. It also explains why agents accept these demands *and* accept criticism when they fall short of these demands (absent the possession of some excuse). And our account reflects the context- and interest-sensitivity of these norms, suggesting that what counts as a reasonable demand varies greatly across a variety of contexts and between various individuals.

Putting all the pieces together, Randy is responsible for his omission in *Bourbon* because he possesses certain capacities that enable him to get the bourbon—capacities that make a general interpersonal bourbon-getting demand reasonable in his particular case, given his saying he would do so—and he was acting within a domain of competence, thereby making the particular demand reasonable. Insofar as Randy is a responsible agent moving within a sphere of competence while he unwittingly omits, he is responsible for failing to get the bourbon.

Briefly, we want to explain the generalizability of our account. According to many different theories of omissions, an agent does not omit to do something unless there is some operative norm (understood broadly) that calls for the agent to do (or avoid doing)

¹⁵ While our discussion of competence is indebted to Raz's (2011) account, we depart from him in significant ways. One notable way is that Raz's account of competence focuses entirely on the agent's *beliefs* about her competence and self-conception. Some have criticized Raz's account on just this point (see, e.g., Watson, 2016) since it seems that agents can be and are deceived (perhaps systematically) about the abilities that they have and the degree to which agents can reliably exercise these abilities (see Kruger & Dunning, 1999). Our account, however, avoids this criticism. While we incorporate agential beliefs into our story, we think that *de facto* reliability is more fundamental. While beliefs about one's abilities and competence has some downward influence on *de facto* reliability, empirical research suggests that beliefs cannot account for the full range of facts about agential competence and reliability (see Protzko & Aronson, 2016; Flore & Wicherts, 2015). In this way, we take it that our account of competence is better suited to the empirical facts than Raz's account.

that thing (cf. Murray, 2018; McGrath, 2005; Clarke, 2014; Henne, Pinillos, and De Brigard, 2017). Thus, any instance of an omission entails norm violation. Part of the content of these norms will include various demands about the allocation of attentional resources (i.e., they demand a certain level of vigilance). By definition, an unwitting omission occurs without the agent's being aware that she will omit or is omitting to behave in some way. So, for any unwitting omission, there will be a lack of awareness (failure of vigilance) that results in (or constitutes) the violation of a norm (part of the content of which is some reasonable demand for a level of vigilance). These are features of unwitting omissions cases generally, so the core aspects of our account will hold for any instance of an unwitting omission.

Conclusion

We have argued that Randy is criticizable for his unwitting omission in *Bourbon* in virtue of failing to be appropriately vigilant. Failures of appropriate vigilance are those where an agent manifests substandard vigilance, where standards of vigilance are determined by the content of reasonable demands made on our cognitive resources. We have also argued that the content of these demands are binding on us in virtue of wanting to manifest reliability within certain domains of competence.

We take our account to be a descriptive project or, insofar as it is normative, the account is independent of any substantive, first-order normative commitments. To be sure, there are complex questions lurking here about why characteristically moralizing responses—e.g., blame, condemnation, sanction—are appropriate or licensed. On these matters, there are a variety of accounts of moral responsibility that strike us as broadly compatible with the kinds of considerations advanced here (see Darwall, 2006; Vargas,

2013-a; McKenna, 2012; Nelkin, 2011; McGeer, 2015; Fricker, 2016). We think our descriptive project can be extended in various ways based on overarching views about the foundations of moral responsibility.

REFERENCES

- Agule, C.K. 2016. "Resisting Tracing's Siren Song," *Journal of Ethics and Social Philosophy* 10:1, 1-24.
- Ahrens, M. (2016). *Home fires involving cooking equipment* (Quincy, MA: NFPA Fire Analysis and Research).
- Amaya, S. 2013. "Slips," In *Noûs* 47:3, 559-76.
- Amaya, S. and Doris, J. 2014. "No Excuses: Performance Mistakes in Morality" *Handbook of Neuroethics*, In J. Clausen and Neil Levy, eds., (Dordrecht: Springer): 253-272.
- Björnsson, G. 2017. "Explaining away epistemic skepticism about culpability," *Oxford Studies in Agency and Responsibility*, In David Shoemaker, ed., Vol. 4 (NY: OUP).
- Brown, Joshua W. 2017. "Models of Anterior Cingulate Cortex Function in Cognitive Control," *The Wiley Handbook of Cognitive Control*, ed. Tobias Egner (London: Wiley), 259-73.
- Buckwalter, W. and Turri, J. 2015. "Inability and obligation in moral judgment," *PLoS One* 10:8, e0136589.
- Clarke, R. 2014. *Omissions: Responsibility, Agency, and Metaphysics* (NY: OUP).
- Cohen, J.D. 2017. "Cognitive Control: Core Constructs and Current Considerations," *Wiley Handbook of Cognitive Control*, In Tobias Egner, ed. (NY: Wiley), 3-28.
- Cohen, J., Dunbar, K. & McClelland, J.L. 1990. On the control of automatic process: A parallel distributed processing account of the Stroop effect," *Psychological Review* 97, 332-61.
- Crocker, J. and Wolfe, C.T. 2001. "Contingencies of Self-Worth," *Psychological Review* 108:3, 593-623.
- Darwall, S. 2006. *The Second-Person Standpoint: Morality, Respect, and Accountability* (Cambridge, MA: Harvard University Press).
- Doris, J.M. 2018. "Collaborating agents: Values, sociality, and moral responsibility," *Behavioral and Brain Sciences* 41, e65.
- Einstein, G.O. and McDaniel, M.A. 2005. "Prospective Memory: Multiple Retrieval Processes," *Current Directions in Psychological Science* 14:6, 286-90.
- Ekstrom, L. 2000. *Free Will: A Contemporary Study* (Boulder, CO: Westview Press).
- Evans, J.S.B.T. and Stanovich, K. 2013. "Dual-Process Theories of Higher Cognition: Advancing the Debate," *Perspectives on Psychological Science* 8, 223-241
- 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.A. 2009. "The Truth about Tracing," *Noûs* 43:3, 531-56.
- Flore, P.C. and Wicherts, J.M. 2015. "Does stereotype threat influence performance of girls in stereotyped domains?" *Journal of School Psychology* 53:1, 25-44.
- Fricker, M. 2016. "What's the Point of Blame? A Paradigm Based Explanation," *Noûs* 50:1, 165-83.
- Gadziola, M.A. and Wesson, D.W. 2016. "The Neural Representation of Goal-Directed Actions and Outcomes in the Ventral Striatum's Olfactory Tubercle," *The Journal of Neuroscience* 36:2, 548-60.

- Graham, P.A. 2014. "A Sketch of a Theory of Moral Blameworthiness," *Philosophy and Phenomenological Research* 88:2, 388-409.
- Guerrero, A.A. 2007. "Don't Know, Don't Kill: Moral Ignorance, Culpability, and Caution," *Philosophical Studies* 136, 59-97.
- Harman, E. 2011. "Does Moral Ignorance Exculpate?" *Ratio* 24:4, 443-68.
- Hartman, R.J. 2016. "Against luck-free moral responsibility," *Philosophical Studies* 173:10, 2845-2865.
- Heine, S.J., Lehman, D.R., Markus, H.R., and Kitayama, S. 1999. "Is There a Universal Need for Positive Self-Regard?" *Psychological Review* 106:4, 766-94.
- Henne, P., Chituc, V., De Brigard, F., and Sinnott-Armstrong, W. 2016. "An empirical refutation of 'ought' implies 'can'," *Analysis* 76:3, 283-90.
- Henne, P., Pinillos, A., and De Brigard, F. 2017. "Cause by Omission and Norm: Not Watering Plants," *Australasian Journal of Philosophy* 95:2, 270-83.
- Higgins, E.T. 1997. "Beyond Pleasure and Pain," In *American Psychologist* 52, 1280.
- Kool, W. and Botvinick, M.M. 2014. "A Labor/Leisure Tradeoff in Cognitive Control," *Journal of Experimental Psychology: General* 143:1, 131-141.
- Kool, W., Gershman, S., and Cushman, F. 2017. "Cost-Benefit Arbitration Between Multiple Reinforcement-Learning Systems," *Psychological Science* 28:9, 1321-1333.
- Kruger, Justin and David Dunning. 1999. "Unskilled and Unaware of It: How Difficulties in Recognizing One's Own Incompetence Lead to Inflated Self-Assessments," *Journal of Personality and Social Psychology* 77:6, 1121-1134.
- Mason, E. 2015. "Moral ignorance and blameworthiness," *Philosophical Studies* 172, 3037-3057.
- McGeer, V. 2015. "Building a Better Theory of Responsibility," *Philosophical Studies* 172:10, 2635-49.
- McGrath, S. 2005. "Causation by Omission: A Dilemma," *Philosophical Studies* 123:1/2, 125-148.
- McKenna, M. 2012. *Conversation and Responsibility* (NY: OUP).
- McKenna, M. 2013. "Reasons-Responsiveness, Agents, and Mechanisms" *Oxford Studies in Agency and Responsibility*, In David Shoemaker, ed., Vol. 1 (NY: OUP): 151-83.
- Mele, A.R. 2006. "Practical Mistakes and Intentional Actions," *American Philosophical Quarterly* 43:3, 249-60.
- Miller, D.T., and Monin, B.. 2016. "Moral Opportunities versus Moral Tests," *The Social Psychology of Morality*, In J. Forgas, L. Jussim, and P. van Lange (NY: Psychology Press), 40-55.
- Murray, S. 2017. "Responsibility and Vigilance," *Philosophical Studies* 174:2, 507-27.
- Murray, S. 2018. "Reference Fiction, and Omission," *Synthese* 195:1, 235-57.
- Murray, S., Murray, E., Stewart, G., Sinnott-Armstrong, W., and De Brigard, F. Forthcoming. "Responsibility for Forgetting," *Philosophical Studies*.
- Nelkin, Dana Kay. 2011. *Making Sense of Freedom and Responsibility* (Oxford: OUP).
- Posner, M.I., and Snyder, C.R.R. 1975. "Attention and cognitive control" *Information processing and cognition: The Loyola Symposium*, In R.L. Solso, ed. (Hillsdale, NJ: Erlbaum Associates).
- Protzko, J. and Aronson, J. 2016. "Context Moderates Affirmation Effects on the Ethnic Achievement Gap," *Social Psychological and Personality Science* 7:6, 500-507.
- Raz, J. 2011. *From Normativity to Responsibility* (Oxford: Oxford University Press).
- Rosen, G. 2008. "Kleinbart the Oblivious and Other Tales of Ignorance and Responsibility," *Journal of Philosophy* 105:10, 591-610.
- Roskies, A.L. 2010. "Don't Panic: Self-Authorship without Obscure Metaphysics," *Philosophical Perspectives* 26, 323-42.
- Roskies, A.L. 2012. "How does the neuroscience of decision making bear on our understanding of moral responsibility and free will," *Current Opinion in Neurobiology* 22, 1022-26.
- Schapiro, T. 2001. "Three Conceptions of Action in Moral Theory," *Noûs* 35:1, 93-117.
- Shabo, Seth. 2015. "More Trouble with Tracing," *Erkenntnis* 80, 987-1011.

- Shenhav, A., Botvinick, M.M., and Cohen, J.D. 2013. "The expected value of control: An integrative theory of anterior cingulate cortex function," *Neuron* 79, 217-40.
- Sher, G. 2006. "Out of Control," *Ethics* 116:2, 285-301.
- Sher, G. 2009. *Who Knew: Responsibility without Awareness* (Oxford: Oxford University Press).
- Shoemaker, D. 2015. *Responsibility from the Margins* (Oxford: Oxford University Press).
- Sinnott-Armstrong, W. 1985. "'Ought to Have' and 'Could Have'," *Analysis* 45:1, 44-48.
- Smith, A.M. 2005. "Responsibility for Attitudes: Activity and Passivity in Mental Life," *Ethics* 115:2, 236-71.
- Sripada, C. 2015. "Moral responsibility, reasons, and the self," In D. Shoemaker (ed.), *Oxford studies in agency and responsibility* (vol. 3, pp. 242-264) (Oxford: Oxford University Press).
- Sripada, C. 2016. "Self-expression: A deep self theory of moral responsibility," *Philosophical Studies* 173, 1203-1232.
- Timpe, Kevin. 2011. "Tracing and the Epistemic Condition on Moral Responsibility," *The Modern Schoolman* 88(1-2), 5-28.
- Vargas, M. "Why the Luck Problem Isn't," *Philosophical Issues* 22:1, 419-36.
- Vargas, M. 2013a. *Building Better Beings* (Oxford: Oxford University Press).
- Vargas, M. 2013b. "Situationism and Moral Responsibility: Free Will in Fragments," *Decomposing the Will*, eds. T. Vierkant, J. Kiverstein, and A. Clark (Oxford: Oxford University Press).
- Watson, G. 2016. "Raz on Responsibility," *Criminal Law and Philosophy* 10:3, 395-409.
- Zimmerman, M. 1986. "Negligence and Moral Responsibility," *Noûs* 20, 199-218.