Philosophia (2018) 46:831–855
https://doi.org/10.1007/s11406-018-9968-7

Moral Responsibility, Voluntary Control, and Intentional Action

Kyle G. Fritz[1]

Received: 13 July 2017 / Revised: 29 December 2017 / Accepted: 29 March 2018 /
Published online: 13 April 2018
© Springer Science+Business Media B.V., part of Springer Nature 2018

Abstract Many theorists writing about moral responsibility accept that voluntary control is necessary for responsibility. Call such theorists *volitionists*. Recently, volitionism has been called into question by theorists I call *nonvolitionists*. Yet neither volitionists nor nonvolitionists have carefully articulated a clear volitionist thesis, nor have they sufficiently explained the concept of voluntary control that somehow seems connected to volitionism. I argue that attempts to explain the volitionist thesis, voluntary control, and their relation are more problematic than have previously been recognized. Instead, I recommend understanding volitionism in terms of intentional actions and omissions. This understanding has several benefits. It clarifies the debate and its parameters, it avoids the problematic notion of voluntary control while relying on the clearer notion of intentional action, and it highlights that the debate between volitionists and nonvolitionists essentially concerns the nature and scope of obligations. As a result, understanding volitionism in terms of intentional actions and omissions can help breathe new life into the volitionist debate.

Keywords Moral responsibility · Voluntary control · Intentional action · Omissions · Volitionism · Obligation

Many theorists hold that to be morally responsible for an action or omission, one must have acted or omitted voluntarily. Similarly, to be responsible for some consequence or attitude, one must have brought about that consequence or attitude through one's voluntary behavior. Call such theorists *volitionists*. Recently, some have begun to question volitionism. These theorists, who I will call

✉ Kyle G. Fritz
kgfritz@olemiss.edu

[1] Department of Public Policy Leadership, The University of Mississippi, PO Box 1848, 105 Odom Hall, University 38677-1848 MS, USA

nonvolitionists, cite mundane cases in which agents seem to be responsible for nonvoluntary items. Consider the following:

> *Bad Joke:* Ryland is very self-absorbed. Though not malicious, she is oblivious to the impact that her behavior will have on others. Consequently, she is bewildered and a bit hurt when her rambling anecdote about a childless couple, a handicapped person, and a financial failure is not well received by an audience that includes a childless couple, a handicapped person, and a financial failure. (Sher 2009: 28)

> *Airport:* Josh promised to pick up Kate from the airport at 8:00. Despite his promise and his generally reliable memory, he forgets about picking up Kate and fails to fulfill his promise.[1]

Ryland and Josh both seem responsible for some item that is not voluntary. Although Josh did not choose or intend to forget to pick up Kate and did not voluntarily fail to pick her up, it seems fitting for him to feel guilty for his omission and for Kate to feel some mild resentment toward him for breaking his promise. Similarly, Ryland's audience is justified in being indignant not only at her joke, which was voluntary, but also at her insensitivity, which was not. Ryland should be more aware of the impact her behavior has on others. Nevertheless, she doesn't voluntarily ignore others' welfare; she simply fails to appreciate the impact of her behavior.

These sorts of cases are ordinary. Yet our judgments that agents are responsible for their nonvoluntary faults, including attitudes, seems to be in tension with the commonly held principle that agents can only be responsible for what is voluntary.[2]

Whether voluntary control is required for responsibility seems to lie at the heart of this debate between volitionists and nonvolitionists. Yet volitionists rarely explicitly state a paradigmatic thesis that unites them, let alone argue for such a thesis. On the other hand, nonvolitionists fail to give a sufficiently clear description of the volitionist thesis they reject. This lack of clarity leaves us with only a general idea of the volitionist debate—one which seems to have stalled at least in part due to this lack of conceptual clarity.

Some writers have shifted focus from voluntary control to other features. For instance, Neil Levy, a volitionist, has focused on consciousness instead. Levy argues that consciousness of some action or omission's moral significance is required for responsibility (2014: 87), and standard nonvolitionist views fail to meet this requirement. While a useful way to proceed, this criticism does not clearly get at the heart of the volitionist debate.[3] Yet to find a productive way forward, it would help to have a better idea of the true nature of the debate.

In this paper, I recommend understanding volitionism in terms of *intention* rather than voluntary control. This understanding has several benefits. There are several unappreciated complications with the notion of voluntary control, yet our concept of

[1] For similar "forgetting cases", see Smith 2005: 236 and Clarke 2014: 164.

[2] One might argue that this tension is merely apparent by appealing to a tracing strategy. I discuss that strategy below in 1.1.

[3] Levy's goal is not merely to criticize nonvolitionism, so this is not intended as a criticism of Levy's work or methodology.

intentional action is sufficiently clear. Additionally, this new understanding captures that the debate between volitionists and nonvolitionists essentially concerns the scope and nature of obligations. By shifting our focus to intentional actions and omissions, I argue, volitionists and nonvolitionists can move forward constructively by focusing on this underlying issue of obligation.

In section 1, I begin with some distinctions fundamental to understanding the debate between volitionists and nonvolitionists. In section 2, I search for and offer traditional volitionist theses by surveying typical characterizations of volitionism in the literature. Because such theses depend upon some conception of voluntary control, in section 3 I explore an account offered by Robert Adams. I argue that this account fails and suggest that we should look for a more constructive way forward. In section 4, I argue that it is clearer and more useful to think of volitionism in terms of intentional actions and omissions rather than in terms of voluntary control. Finally, in section 5, I explain some benefits of understanding the debate in this way.

1 Preliminaries

To forestall potential misunderstandings, I start by making an important distinction and clarifying a shared notion of moral responsibility.

1.1 Direct Vs. Derivative Responsibility

One response many volitionists make regarding cases like *Bad Joke* and *Airport* is that agents like Ryland and Josh are only indirectly responsible for their failures. Although Josh does not voluntarily fail to pick up Kate, he might still be responsible for so failing in virtue of the fact that he exercised some voluntary control at some point in the past, and he was aware or should have been aware that the way in which he exercised that voluntary control would lead to him not picking up Kate. For example, perhaps Josh could have written himself a reminder, but chose not to despite being aware that he might forget his promise as a result. Similarly, perhaps Ryland voluntarily behaved in certain ways in the past such that she foresaw or should have foreseen that she would become insensitive.

This tracing technique is a common strategy for those who claim that voluntary control is required for responsibility.[4] We can distinguish between *direct* responsibility and *derivative* responsibility. When *S* is responsible for something, *X*, in virtue of being responsible for something else, *Y*, that does not include *X*, then *S* is *derivatively* responsible for *X*. For example, when Ted drunkenly hits a pedestrian with his car, he is derivatively responsible for hitting the pedestrian. Although Ted does not voluntarily hit the pedestrian, Ted voluntarily drank to excess when he was aware that he would have to drive shortly. Supposing that Ted is responsible for his drinking, he is derivatively responsible for hitting the pedestrian while drunk in virtue of his prior choice to drink. When *S* is responsible for *X* but not solely in virtue of being responsible for any other thing, *Y*, *S* is *directly* responsible for *X*. Ted is directly responsible for his

[4] Clarke 2014 seems skeptical of tracing's ability to explain responsibility for certain omissions, and Vargas 2005 raises general worries for tracing. Fischer and Tognazzini 2009 defend tracing against Vargas's worries.

choice to drink, as his responsibility for this choice does not depend solely on his responsibility for something else.

Volitionists and nonvolitionists debate about the necessary conditions of *direct* responsibility. Volitionists, unlike nonvolitionists, argue that voluntary control is necessary for direct responsibility. The parties disagree about how best to explain our responsibility for items that we did not voluntarily control. While many volitionists might trace responsibility for these items to some previous exercise of voluntary control, nonvolitionists will often hold agents directly responsible for failures to notice or for having or lacking certain attitudes.[5]

1.2 Kinds of Responsibility

Some have claimed that the debate between volitionists and nonvolitionists is not substantive, as volitionists appeal to one kind of responsibility and nonvolitionists to another.[6] If so, volitionists and nonvolitionists may be talking past one another.

One purported kind of responsibility is *attributability*. According to Gary Watson, an agent is attributability-responsible for whatever thoughts or actions express her practical identity or commitments (2004). Attributability is often contrasted with *accountability*, which concerns demands and obligations such that when an agent has violated some such demand, she can rightly be subject to the reactive attitudes (e.g., resentment, indignation, or guilt) or sanctions. As Clarke, McKenna, and Smith point out, Watson's understanding of accountability raises issues about fairness and our authority to impose sanctions that mere attributability apparently does not (2015: 4).

Some reject this distinction, arguing that responsibility should be understood as a unified concept of *answerability*. An agent is answerable for some item if we can sensibly demand from her the reasons that she took to justify that item (Smith 2008, 2012; Hieronymi 2008). According to Smith, these reasons can be good or bad, and an agent may be subject to moral appraisal—including the reactive attitudes—depending on these reasons.

I will avoid this debate and focus on a shared notion of moral responsibility that already seems to exist in the literature. Both volitionists and nonvolitionists are interested in whether the reactive attitudes are justified or deserved with regard to some item. Levy, who identifies as a volitionist, writes, "To say that an agent is morally responsible (for an act, omission or attitude) is to say that the Strawsonian reactive attitudes are justified in relation to her with regard to that act, omission or attitude" (2005: 1). Angela Smith, a nonvolitionist, writes something slightly different: "to say that someone is 'morally responsible' for some thing is to say that she is open, in principle, to moral appraisal—including moral praise and blame—on the basis of it" (2008: 370). But she is quick to clarify that she only avoids talk of the reactive attitudes

[5] Some volitionists might claim that we are neither derivatively nor directly responsible for many or any of these nonvoluntary items. If Ryland is not voluntarily insensitive, and if she is not derivatively responsible for her insensitivity in virtue of her past voluntary behavior, then she is not responsible at all for the insensitivity. Nevertheless, this insensitivity is still bad (see Levy 2005). Notice, however, that if tracing fails often for a number of nonvoluntary items for which we typically hold agents responsible, volitionists who favor tracing must accept that we are responsible for much less than we originally thought.

[6] See Levy and McKenna 2009: 118. The metaphysical status of this distinction in responsibility is unclear, as writers refer to *kinds*, *senses*, or *faces*. I intend to be noncommittal on this issue but use "kind" for consistency.

to make room for morally neutral acts, which may not warrant any reactive attitude, and because "whether a Strawsonian reactive attitude is 'justified' in a particular context might involve considerations that have nothing to do with the agent's responsibility and blameworthiness" (370, n. 9).[7]

The kind of responsibility I take to be at issue in the debate between volitionists and nonvolitionists is whatever kind warrants the reactive attitudes.[8] When I say that an agent is responsible for something, I mean that the agent may be worthy of some reactive attitude on the basis of that thing.[9] Here I will be concerned largely with negative reactive attitudes, which I understand to be attitudes some agent has in response to the belief that another agent has done wrong, has done something bad, or is morally bad in some way. The paradigm negative reactive attitudes are indignation, resentment, and guilt, though certain types of moral anger may also be negative reactive attitudes as I understand them. Given my characterization of the negative reactive attitudes, one can be indignant not only at another's wrong action, but also at another's objectionable attitudes. Volitionists and nonvolitionists may disagree about whether one's indignation in each case is justified.

2 Traditional Volitionism

Theorists have traditionally characterized volitionism in the following ways:

> *C1*: "[W]e are ethically accountable only for our voluntary actions and omissions." (Adams 1985: 3)
>
> *C2*: "[C]hoice or voluntary control is a precondition of legitimate moral assessment." (Smith 2005: 236)
>
> *C3*: "[We can distinguish] two opposed accounts of moral responsibility, the account sometimes called attributionism, which stresses the extent to which an action is reflective of an agent's real self, and volitionism, which stresses choice and control." (Levy 2008: 214)

C1-C3 emphasize that choice and voluntary control are required for responsibility. Yet one might naturally wonder what is meant by *voluntary control*, and how choice is related to it. Unfortunately, most theorists have either been entirely silent with respect to these questions or else have merely gestured at answers.

[7] Here Smith references considerations in the ethics of blame, such as whether the blamer has the moral standing to blame (see Smith 2007 and Coates and Tognazzini 2012).

[8] I set aside sanctions, which unlike the reactive attitudes, aim to cause suffering in another. Since the reactive attitudes may be felt and expressed without any such intention, they are importantly different from sanctions. Perhaps one can understand the sort of responsibility with which I am concerned as accountability, given accountability's connection with the reactive attitudes. Some might protest that accountability concerns obligations, and one cannot be obligated to have attitudes, so accountability cannot be the kind of responsibility at issue. I return to this point below in section 5.

[9] I will leave aside the issue of morally neutral actions here that concerned Smith 2008 above. I will also focus on whether an agent is *blameworthy* for some item, but not on whether the agent *should be blamed* for the item.

Michael McKenna has helpfully illustrated the lack of clarity that writers use in describing volitionism. He notes that Angela Smith characterizes volitionism as restricting responsibility to “what an agent has *voluntary control* over,” or “whatever an agent has *direct control* over,” or “what an agent has *chosen*,” or “what an agent has *voluntarily chosen*,” or what an agent has *deliberately chosen* (McKenna 2008: 30). This lack of clarity is alarming, because these different characterizations of volitionism may capture different views, depending on one’s theories about control and voluntariness.

To be fair, volitionism is a family of views rather than one particular view. What appears to be a lack of clarity may simply be Smith’s attempt at inclusiveness: “What I have called the volitional view of responsibility is actually better understood as a cluster of distinct views which share a common assumption, namely, that choice, decision, or susceptibility to voluntary control is a necessary condition of responsibility (for attitudes as well as actions)” (Smith 2005: 238).

Smith identifies several types of volitionist views that are united by this assumption. Two are of interest to me here.[10] *Exercised capacity views* claim that we must be able to trace some action, omission, attitude, or consequence to a previous exercise of voluntary control in order for the agent to be responsible for that item. Though Smith calls such views “prior choice” views, I prefer my terminology for multiple reasons. While choosing is an exercise of voluntary control, targeting only “a person’s own prior choices or decisions” (239) as anchor points for tracing rules out omissions. Additionally, as I will show below, many of the actions for which we seem directly responsible do not trace back to choices at all, and volitionists might object to such a narrow conception of their view. Exercised capacity views ground direct responsibility in the exercise of voluntary control (sometimes via choice, though not necessarily), and responsibility for any nonvoluntary item such as attitudes or forgettings must be traced back to some previous exercise of voluntary control, as I described in 1.1 above. On what I will call *capacity* views, direct responsibility is not rooted in exercising some capacity (such as making an actual choice), but rather simply in one’s *capacity* to voluntarily control the item.

The difference between exercised capacity views and capacity views is significant. Despite these differences, both exercised capacity and capacity views take “the voluntary to be located in choice,” understanding choice “as involving deliberate undertakings” (McKenna 2008: 30). Given this understanding, and given the differences in the views, it is helpful to begin with two separate volitionist theses:

> *VT1e: S* is directly responsible for *x* only if *S* exercised direct voluntary control over *x*.[11]

[10] Because she focuses on responsibility for attitudes, Smith describes volitionist approaches to responsibility for attitudes. I am concerned with more than attitudes, so Smith’s characterization of different volitionist views does not match my own, but the spirit is the same. I set aside Smith’s endorsement views, as what I say regarding exercised capacity views will likely also apply to endorsement views. Additionally, I focus on capacity views rather than Smith’s voluntary control views because I am interested in retrospective responsibility.

[11] I use the variable *x* here and throughout to stand for actions, omissions, attitudes, and consequences.

> *VT1c: S* is directly responsible for *x* only if *x* was susceptible to *S*'s direct voluntary control.[12]

Advocates of exercised capacity views can endorse VT1e. When *S* chooses to perform some action, *S* exercises direct voluntary control over that choice (and perhaps that action) and so can be responsible for it.[13] Nonvoluntary attitudes and omissions, however, are not items over which we exercise direct voluntary control. Consequently, *S* cannot be directly responsible for these items on VT1e.

Capacity views do not require that any capacity be exercised for an agent to be directly responsible for some item. VT1c reflects this important difference. For instance, if *S* could have chosen to *x* and could have brought *x* about by so choosing, then *S* can still be morally responsible for *x*—even if *S* never made any actual choice to *x*.

One might wonder why I have offered two distinct volitionist theses rather than one unified volitionist thesis that separates volitionists from nonvolitionists. The reason is that a unified thesis will be merely a disjunction of VT1e and c:

> *VT1u*: *S* is directly responsible for some item, *x*, only if *S* exercised direct voluntary control over *x* or *x* was susceptible to *S*'s direct voluntary control.

VT1u provides merely superficial unity. Exercised capacity views and capacity views are united by a focus on the importance of voluntary control to responsibility, but each requires that control to be involved in a different way. That difference in scope prevents any more significant unification. Exercised capacity views are more restrictive than capacity views, so exercised capacity theorists will see the second disjunct as unnecessary. But capacity theorists will read the "or" of VT1u as inclusive and think of the first disjunct as superfluous, given that exercising direct voluntary control over *x* requires that *x* be susceptible to one's direct voluntary control.

A broad thesis like VT1u can separate volitionists from nonvolitionists but fails to respect important differences between volitionist views. VT1e and c make the differences between volitionist views clearer, though VT1e does not distinguish *all* volitionists from nonvolitionists. Which thesis one favors depends on one's goals. Regardless, any version of VT1 remains obscure without some account of voluntary control.

[12] Note that VT1c includes all the items VT1e does, because if *S* exercised direct voluntary control over *x*, then *x* was susceptible to *S*'s direct voluntary control.

[13] The way I have characterized direct and derivative responsibility leaves open that one might be both directly and derivatively responsible for some item, such as one's voluntary actions. Alternatively, one might think that agents can only exercise direct voluntary control over their choices, not their chosen actions, so perhaps direct voluntary control is limited to choices. For more on direct control, see Mele 2017. Here I leave it open that agents might exercise direct voluntary control over both the choice to *A* as well as over *A*-ing itself.

3 An Account of Voluntary Control

Robert Adams, one of the first theorists to seriously question volitionism, understands voluntary control at least partially in terms of choice. He offers the following account of direct voluntary control:

> *VCd*: "To say that something is (directly) within my voluntary control is to say that I would do it (right away) if and only if I (fully) tried or chose or meant to do so, and hence that if I did it I would do it *because* I tried or chose or meant to do it, and in that sense voluntarily" (1985: 8).

As Adams writes, "the key to this conception [of voluntary control] is the idea of a trying, choosing, or meaning of which what is voluntarily done is (or would be) the object" (8).[14] He focuses on choosing at least in part because "in a voluntary action the action itself is an object aimed at by the agent in a way that is aptly described as 'trying, choosing, or meaning'" (8).

At first glance, VCd might appear to properly classify relevant items as within one's direct voluntary control. Suppose Sam raises her arm. According to VCd, to say that Sam's arm raising is within her direct voluntary control is to say that she would raise her arm just in case she fully tried or chose or intended to raise her arm, and she would raise her arm because of this trying, choosing, or intending. If some manipulation of Sam's brain during surgery causes her arm to rise, we would not count this as voluntary. Appropriately, VCd does not classify such arm raising as an exercise of voluntary control either.

Similarly, VCd illustrates how one's attitudes are not within one's direct voluntary control. One cannot choose or intend or try to become angry and thereby make oneself angry right away, but can only engage in certain behavior to make oneself angry. Suppose that Val decides to dwell on politics, knowing that it will make her angry. In a sense, Val's becoming angry is voluntary, but it is not directly so, as it is not immediate. Her anger is only within her *indirect* voluntary control.

Adams offers a conception of indirect voluntary control as well:

> *VCi*: "Doing *x* is not directly but only indirectly within my voluntary control if I could try successfully to do *x* right away, but only by trying successfully to do something else as a means to doing *x*" (8, n. 5).

Volitionists hold that agents are only derivatively responsible, if responsible at all, for their attitudes. Thus the focus for nonvolitionists is on direct voluntary control as spelled out in VCd rather than the extended sense of "voluntary" in VCi.[15,16] Because

[14] I will understand "meaning" and "intending" to be equivalent.

[15] Similarly, Pamela Hieronymi notes that if one includes in the class of 'the voluntary' "not only voluntary activities but also those things that are the (possible) result of such activity" (2008: 358, n. 1), believing can be the result of voluntary activity. But those who debate the voluntariness of believing are not thinking of this extended sense.

[16] VCi is too weak. Recall Ted, the drunk driver. He exercised his direct voluntary control in drinking alcohol. But Ted did not drink and then drive as a means to successfully hit a pedestrian. Nevertheless, it seems that Ted's hitting the pedestrian was within his indirect voluntary control.

volitionists are focused on direct voluntary control, and because attitudes are not under one's direct voluntary control according to VCd, attitudes are not something for which agents are directly responsible according to volitionists. Additionally, what agents forget or neglect is only under their indirect voluntary control.

Nevertheless, VCd is problematic as an account of direct voluntary control, as it falls prey to various sorts of counterexamples. Suppose one chooses to *A* and *A*s because one so chooses, but would have *A*-ed even if one hadn't chosen, tried, or intended to *A*. For instance, Trina might choose to remain sitting at her desk and remain there as a result of her choice. But if she hadn't chosen to remain at her desk, perhaps she would have continued deliberating and would have remained at her desk without trying or intending to do so as she deliberated. Such cases create problems for the necessary condition of the biconditional in VCd.[17]

Relatedly, VCd appears to fall prey to similar sorts of counterexamples as the classical compatibilist's conditional analysis of "could have done otherwise".[18] Suppose Abdul suffers a phobia of spiders. He cannot even try, choose, or intend to pick up a friend's pet tarantula. It doesn't seem that picking up the tarantula is within his voluntary control, yet according to VCd it is. Additionally, several sorts of intentional actions require significant skill to perform. Not even a professional basketball player can sink a free throw on every attempt, even though she might try, choose, or intend to do so. VCd cannot obviously accommodate skilled intentional actions.

Beyond these counterexamples is yet a further problem for VCd. Volitionists do not only want to capture intentional overt actions as voluntary, but also a variety of mental actions, including choosing and willing. Choosing and willing seem to be within one's direct voluntary control. If we understand direct voluntary control as in VCd, however, we must say that *S*'s choosing to *x* is within her direct voluntary control just in case *S* would choose to *x* iff *S* tried or chose or intended to choose to *x*. This is problematic for two reasons. First, it is often false that in choosing to *x*, *S* chooses to *x* because *S* chose to choose to *x*. This would lead to an unwelcome regress. More importantly, however, we do not think that choosing is under one's voluntary control just in case one chooses to choose.

While Adams does not acknowledge this issue, he seems aware that VCd will have problems with choosing, trying, and intending. He writes:

> Trying, choosing, or meaning can itself be regarded as voluntary, as an operation of the will, so to speak, that is characteristic of voluntary action. It may perhaps also be said, in an extended sense, to be within our voluntary control, though it is itself the controlling, rather than an object of control. (9)

This extension of VCd to cover trying, choosing, and intending will not do, however. First, the extension seems ad hoc. Adams began with things that one does as a result of choosing and then attempted to extend the concept of direct voluntary control backwards to cover choosing as well, despite choosing not satisfying the original parameters of his account.[19]

[17] I owe this counterexample to Randy Clarke.

[18] For more detail see van Inwagen 1983: 114–121.

[19] It seems that any conception of voluntary control that begins with such items will fail to capture vital mental actions as within our voluntary control. Hieronymi's account of the voluntary also seems to rule out choice for similar reasons (2008: 366).

Second, it is not clear that trying, choosing, and intending should all be regarded as voluntary. As Holly Smith writes, "[m]any people believe that desires and intentions are not typically themselves voluntary, but ultimately arise from our biology, our environments, and our prior experiences" (2011: 123). Additionally, one can non-actively come to intend to *A*. When Al unlocks his office door in the morning, he intends to unlock his door. But he does not actively acquire this intention. He is not uncertain about whether to unlock the door, so there is no need for him to actively form the intention to unlock the door by choosing to unlock it (Mele 2003: 200–201). In coming to intend to unlock the door, Al has not voluntarily done anything. If this is right, then not all of these items should be understood as being within our direct voluntary control. Thus Adams cannot simply extend VCd to include choosing, intending, or trying. Non-actively acquired intentions are not voluntary, though the actions that issue from these intentions may be voluntary.

To summarize, VCd cannot be an accurate conception of direct voluntary control. Besides being susceptible to numerous counterexamples, choosing, a paradigm action for which we are directly responsible and which is an exercise of our voluntary control, is not captured on VCd. But any extension of VCd to capture choosing, intending, or trying is ad hoc at best, and such an extension will require further refinements to rule out non-actively acquired intentions that should not count as voluntary.

Are there ways to refine VCd or proceed with something similar in spirit to get the right result? One reason why VCd fails is because it cannot easily accommodate choosing itself. The things we do as a direct result of choosing are directly voluntary, but so are our choosings. *S* does something as a *direct* result of choosing to *x* when *S*'s choosing to *x* brings about *x* without needing to bring about some intermediary thing first. *S* need not choose to *y* in order to bring about *x*. Learning from Adams's mistake, we might better understand direct voluntary control by starting not with the things one does as a direct result of choosing, but with choice and these mental actions that are voluntary themselves.

Choosing is not the only mental action that is voluntary, however. We can try to do *x*, will to do *x*, and intend to do *x*—all without choosing to do *x*. In fact, most of our everyday behavior is unthinking and habitual, involving no choice at all, but it is no less voluntary for that. Perhaps then, we should focus on *acts of will* as voluntary, such that acts of will and the things we do as a direct result of some act of will are directly voluntary.[20] Unfortunately, it is unclear precisely what besides choosing counts as an act of will.

Adams mentions trying and intending as voluntary. Given that trying to *A* is initiated by some intention, it seems that trying should be understood as an act of will. With Adams and Mele, I hold that "trying to *A* is an event or process that has *A*-ing as a goal and is initiated and (normally) sustained by a pertinent intention" (1992: 326). So trying is involved in all intentional action, and successful tryings to *A* are *A*-ings.

Yet what portion of tryings should count as acts of will? Randolph Clarke and Thomas Reed argue that only beginning portions of such attempts should count, when one's intention activates certain motor neurons. The entire trying is not an act of will. When a golfer moves her hands to sink a putt, such movements count as part of the attempt, yet not as part of her act of will. Hand movements, unlike a portion of the

[20] I understand acts of will and willings to be equivalent but use "acts of will" for stylistic reasons.

willing, may be slow (Clarke and Reed 2015: 10–11). On the other hand, Brian O'Shaughnessy 2008 insists such hand movements are part of the act of will. Although it seems clear that at least some portions of tryings should be understood as acts of will, it is less clear *what portion* should count. What counts as an act of will may impact how we understand what *S* brings about as a direct result of some act of will.

As I mentioned above, some have argued that intending is no more voluntary than believing (Hieronymi 2008: 368). If so, intending should not count as an act of will. Nevertheless, one might think that *forming* an intention counts. Sometimes one comes to have an intention to *x* by deciding to *x*, and deciding is an intentional action (Mele 2003: 186–187). But if choosing and deciding are equivalent, then we have already included such intention formation above.

Answering these questions does not settle the matter, as there may be further candidate acts of will. One might insist that volitions, whatever they may be, are acts of will (Zimmerman 1988; Ginet 1990). Or one might insist, like O'Shaughnessy 2009, that one may only try when one may fail, and certain sorts of mental actions like imaginings do not include this opportunity for failure. If so, then the way that I have described tryings above is incorrect, and there is some further sort of act of will. Settling the matter may require splintering volitionists even further depending on their favored understanding of what counts as an act of will.

These are challenging questions within the philosophy of action. While there may be some way to answer them, the volitionist debate will only stall further if it is bogged down with such matters. Perhaps we should look for a different way forward.

4 Focusing on Intentional Actions and Omissions

McKenna observes that Smith and other theorists understand "the voluntary to be located in choice," and they conceive of choice "as involving deliberate undertakings" (2008: 30). While choice involves a deliberate undertaking, we have seen that it is too narrow to capture everything a volitionist plausibly wants to capture. One reason for this may be that volitionists and nonvolitionists alike are using "choice" and "choosing" in a broader sense than many action theorists. Instead, they seem concerned with deliberate undertakings in general. Acts of will such as choosings are intentional, deliberate undertakings. If this is right, rather than focus on *choice* and *voluntariness*, we might make more progress by focusing instead on *intentional actions and omissions.*

Such a move is not unmotivated. In an argument about epistemic justification, William Alston understands voluntary control "on the model of intentional action" (1988: 277). He distinguishes between basic voluntary control, non-basic immediate voluntary control, and long-range voluntary control. Only the first two concern me, because only they fall under the term "direct control" as Alston conceives it (269). Basic voluntary control is exercised in performing "'basic actions', actions we perform 'at will', just by an intention, volition, choice, or decision to do so" (260). Non-basic immediate voluntary control involves the agent being "able to carry out the intention 'right away', in one uninterrupted intentional act, without having to return to the attempt a number of times after having been occupied with other matters" (269).

While there are familiar pitfalls here, Alston's shift from voluntary control to intentional action is important. In a response to Alston's argument, Chuard and Southwood summarize Alston's understanding of voluntary control as requiring (1) the capacity to form intentions and (2) the capacity to carry out such intentions (2009: 603). Rather than understanding voluntary control in terms of choosing, it may be more helpful to reexamine the volitionist debate framed in terms of intentional actions and omissions instead.

I submit that we have a clearer picture of intentional actions and omissions than we do of voluntary control, in part thanks to an account offered by Mele and Moser 1994. Mele and Moser give necessary and sufficient conditions, but two of these conditions are designed to handle cases of lucky action.[21] I am not concerned with such cases here, so I focus only on two necessary conditions that are distinctive of intentional action: "Necessarily, an agent, *S*, *intentionally* performs an action, *A*, at a time, *t*, [only if] (i) at *t*, *S A*-s and her *A*-ing is an action, [and] (ii) at *t*, *S* suitably follows—hence, is suitably guided by—an intention-embedded plan, *P*, of hers in *A*-ing" (Mele and Moser 1994: 63).

The first condition is rather trivial, but the second warrants further explanation. In performing an intentional action, one has an *aim* to perform an action of some type. Mele and Moser note that "Our performing an action of type *A* intentionally requires our having some representation of that action type, or at least some representation that can suitably guide our performing an action of that type" (44). This representation, as Mele and Moser understand it, is an *action plan*, which may be relatively simple (a representation of one's *A*-ing) or more complex (a representation of a plan including the means by which one will achieve some goal) (42). This representation must guide the action in order for the action to be intentional.

So much for intentional action. What about omissions? Although philosophers disagree about the nature of omissions, Clarke argues convincingly that many omissions are simply absences of action (2014: chs. 1–2). Some omissions may be negative actions (e.g., doing one thing to prevent oneself from doing another), while others may be positive actions linguistically described in a negative way (e.g., holding still). But not all omissions can be accounted for in these ways, so omissions cannot simply be treated as a special kind of action.

Focusing only on intentional actions in the volitionist debate, then, would get at only part of the story. We can omit to act, and such an omission can be intentional or not. Volitionists and nonvolitionists alike are concerned with explaining how agents are responsible for omissions. Although we may explain responsibility for some omissions derivatively, we may be directly responsible for others.

I suggested above that there is no illuminating way to express one simple volitionist thesis that accurately captures all and only volitionists. While all volitionists insist that voluntary control is required for direct responsibility, they differ regarding how voluntary control is required. At best, we can offer a merely disjunctive thesis to provide some unity between exercised capacity theorists and capacity theorists. This same problem arises once again when we think of volitionism in terms of intentional actions

[21] These are conditions (iii) and (iv) (1994: 63).

and omissions. So I offer two separate theses again, which roughly parallel exercised capacity views and capacity views, respectively:

> *VT2e: S* is directly responsible for some item, *x*, only if *x* is some intentional action or intentional omission of *S*'s.

> *VT2c: S* is directly responsible for some item, *x*, only if *x* is some intentional action or omission of intentional action of *S*'s.

Before looking at their differences, I'll begin with their similarities.

Like VT1e and c, VT2e and c entail that agents can be directly responsible for many items that volitionists wish to capture under the scope of direct responsibility, including mental actions like choosing or deciding.[22] This is vital; just as decisions were not clearly within one's voluntary control on Adams's account, if decisions are not clearly intentional actions, then the motivation for this shift to intentional action is severely undercut.[23] Thus it is worth explaining how choosing or deciding is an intentional action rather than, say, nonactional.

First, I should reiterate that I am only concerned with practical deciding (deciding to act) as opposed to cognitive deciding (deciding that something is the case). Perhaps cognitive deciding is nonactional, such that one simply acquires a belief that *p* is true on the basis of reflection (Mele 2003: 198). Yet even if cognitive deciding is not an action (and so not an intentional action), I favor the widely held view that practical deciding is a momentary mental action of intention formation brought about when an agent is uncertain about what to do (197–202).[24] More specifically, I favor Mele's view that practical deciding is a mental action in which one assents to some first-person action plan: "In deciding to act, one forms an intention to act, and in so doing one brings it about that one has an intention that incorporates the plan to which one assents" (210).

Decisions are actions because they are not immediately acquired, even on the basis of practical reflection. Some intentions might be nonactionally acquired, such as when some desire or perception issues straightaway in an intention, or when there is no uncertainty about what to do (as in the case of Al intentionally unlocking his door above). But not all intentions are like this. Sometimes an agent must resolve to do *B* despite the urge to do *A*, and this resolving is a mental action (202). Other times, even if an agent judges that it would be best to *A* on the basis of practical reflection, this does not guarantee that the agent will immediately acquire an intention to *A* without some further act. Consider Mele's case of Joe, who judges it best for him to quit smoking starting this evening. Despite this judgment, Joe has not yet decided to quit, as he has not yet formed the intention to quit smoking. This formation of the intention would be an action (199–200).

On Mele's view, decisions are not only mental actions, but *intentional* actions. The intention that plays a role need not be as specific as in other cases of intentional actions. That is, one need not intend to decide to *A* in order to decide to *A*. Instead, one may

[22] As I mentioned above, I understand choosing and deciding to be equivalent.

[23] Thanks to an anonymous referee for raising this concern and pressing me to say more on this point.

[24] Though I focus on Mele 2003's view, many theorists besides Mele endorse this claim, including Kane 1996: 24 and McCann 1998: 163.

simply have the intention to decide what to do, and this intention may produce the decision—that is, the action of forming the intention to *A* (202–205). Thus decisions are momentary mental actions of intention formation that resolve practical uncertainty, and they are caused by intentions to decide what to do.[25] As intentional actions, decisions are clearly captured by VT2e and c.

Additionally, a variety of other intentional actions, such as raising an arm, singing a song, or driving a car might be included within the scope of direct responsibility. Importantly, however, they need not be since both versions of VT2 provide only a necessary condition on direct responsibility. This accommodates volitionists who hold that agents are directly responsible only for basic actions. VT2e and c simply restrict the class of items for which agents can be directly responsible. This class might sensibly be restricted even further depending on one's other commitments.

Both VT2e and c exclude from direct responsibility many items that volitionists want to exclude. Attitudes are not actions, so they are not intentional actions. Nor are they omissions. Consequently, they are not items for which one may be directly responsible. Of course, this is not to say that one cannot be directly responsible for intentional actions or omissions one performs to bring about some attitude.

The important difference between VT2e and c is in which omissions they include. VT2c is wider than VT2e, as it does not require that some omission be intentional. Randy might promise his wife that he will buy milk on the way home from work, yet nevertheless forget to stop on the way home. Randy's omission is an absence of an intentional action. Yet Randy's omission is not an intentional omission; he did not intentionally omit to get the milk. He simply forgot.[26]

VT2e and c throw into sharp relief the difficulty of providing any unifying volitionist thesis. Because VT2c is merely a necessary condition and the broadest thesis, it is technically compatible with all volitionist views. Nonvolitionists would reject it, so it can serve as a way to separate volitionists from nonvolitionists. Nevertheless, notice just how broad VT2c is. It leaves open the possibility that one is directly responsible for unintentional, nonvoluntary omissions. Direct responsibility for these

---

[25] Joshua Shepherd raises an important challenge for understanding practical deciding as an intentional action by highlighting a difference between the intentions involved in overt intentional actions as compared to the intentions involved in decisions. The intentions relevant to overt intentional actions "initiate, sustain, and guide action" (2015: 337). One simply matches one's behavior with the content of the intention. But as Shepherd points out, the intentions relevant to practical decisions do not seem able to play this sort of role, as the intentions involved in decisions are open-ended, lacking specific content. Rather than intending to decide to *A*, one simply intends to decide what to do. But if the intention involved in a decision lacks the specific sort of content that one is supposed to match with one's behavior, it's unclear how the agent can have the relevant guidance and thus control over a specific decision such that we can call it an intentional action (337). In light of Shepherd's challenge, one might think that decisions are not intentional actions after all, but perhaps are nonactional instead.

Shepherd does not specify the sort of control with which he is concerned, but I am happy to once again set aside detailed talk of control and adopt Shepherd's own solution to this problem. Shepherd points out that the momentary intention formation of a decision is "an active expression of an agent's skilled deliberative activity…that results from the causal work of a relevant intention, in conjunction with the agent's attentional attunement to relevant features of the deliberative situation" (348). This, Shepherd says, is in contrast with nonactionally acquired intentions, which are not similarly connected to intentions to decide what to do and need not involve attention. Thus there are important differences between nonactionally acquired intentions and the intention formation of decisions, and these differences are sufficient to show that decisions are indeed intentional actions.

[26] This case comes from Clarke 2014: 164 ff.

sorts of unintentional omissions is one of the chief motivations nonvolitionists offer for their view (see cases in Smith 2005). For example, recall *Airport*. Exercised capacity theorists either deny that Josh is responsible for not picking up Kate, or they claim that he is derivatively responsible in virtue of some previous intentional omission or action (see Levy 2005). But VT2c might allow that Josh is directly responsible for his unintentional omission, as it is an omission of an intentional action. Many volitionists may balk at such claims. Those volitionists who favor more restrictive views might prefer a thesis that more accurately represents their view, such as VT2e.[27] For this reason, I think it is unhelpful to offer just one volitionist thesis or to try to find a parallel VT2u. Of course, this does not mean that VT2e and c have nothing in common. Both theses acknowledge that intention in actions and omissions matters, but they disagree about the way in which intention must come into play in omissions.

Before explaining a few benefits of this new suggestion for how to frame volitionism, I turn briefly to another writer who also attempts to state the volitionist thesis without explicitly appealing to voluntary control. Instead, Michael McKenna focuses on the control one is able to exercise by acting freely:

> *VTM:* “A person is morally responsible only for things that are within the control she is able to exercise by acting freely” (2012: 187).

While there is much to admire about VTM, I do not think it will serve as a clear volitionist thesis. Notice that VTM does not distinguish between direct and derivative responsibility. This obscures the debate between volitionists and nonvolitionists. Exercised capacity theorists agree that one’s attitudes may be within the control one is able to exercise by acting freely. But, of course, nonvolitionists accept that one may be responsible for these items as well, though not *in virtue* of the fact that those items were within the control one was able to exercise by acting freely.

Perhaps VTM might be modified to account for the distinction between direct and derivative responsibility. If so, VTM might be comparable to VT2c provided that McKenna is correct about how VTM handles omissions. McKenna claims that “the same control one exercises by acting freely is the control implicated in omissions. In the case of the relevant class of omissions, one refrains from exercising that control” (2012: 188, n. 11). One might omit from exercising this same type of control when one forgets to pick up a friend at the airport, as promised. I prefer my formulations to VTM because they make no appeal to control. It is reasonable to ask a proponent of VTM what kind of control agents are able to exercise by acting freely. If this control is understood as voluntary control, familiar problems may arise once again.

Allow me to summarize. The volitionist debate has stalled, and progress is hindered in part because we lack a clear account of voluntary control. Volitionists and nonvolitionists alike have failed to provide an account that works, but it seems important for knowing what the debate amounts to and how to proceed fruitfully. Volitionists claim that it’s not fair to hold individuals morally responsible for actions or omissions outside their voluntary control. Nonvolitionists demur, insisting that our

[27] Michael Zimmerman explicitly endorses the view that one is directly responsible only for one’s volitions (1988: 40). While many volitionists may find Zimmerman’s view too restrictive, Zimmerman would likely see VT2c as unhelpful given how broad it is.

everyday practices involve doing just this, and these practices are justified. The debate seems to bottom out in this intuition over fairness.

I don't think that it must, however. To make progress, I've suggested that we tackle the problem in a different way. Rather than think about volitionism in terms of voluntary control, we can do so in terms of intentional actions and omissions, as in VT2e and c. Unlike with voluntary control, we have a clear account of intentional action. There is no concern about using "choice" ambiguously or worrying how choice and voluntary control are related. In addition, the respective scopes of VT2e and c are clear. But more than this, they can help us better understand the debate by focusing on the heart of the dispute and point to where the debate must turn to see a way forward.

5 Upshots

Before turning to that benefit, it is worth noting that volitionists and nonvolitionists alike might be suspicious of VT2e and c because the defining feature of volitionism seemed to be voluntary control or choice, yet VT2e and c excise any mention of either. "Intentional" and "voluntary" are not precisely coextensive, as coerced actions may be intentional but not voluntary. And as an anonymous referee points out, it seems that one could satisfy VT2c with little, if any, voluntary control. By shifting discussion to intentional actions and omissions, are we still discussing *volitionism*?

I admit that VT2e and c do not map precisely on to the theses that have been suggested in the debate thus far. In an obvious sense, then, we are not still discussing volitionism as it has been understood, and that is by design. The traditional debate has been unfruitful, so there must be a shift. I have shifted away from talk of voluntariness and control to talk of intentional action, so if one understands volitionism as essentially requiring talk of voluntary control, then one will be inclined to think this shift is merely a change of topic and therefore reject my proposal out of hand.

I am suggesting that it is not clear that volitionism essentially is or must be a discussion about voluntary control, so this shift is not a complete change of topic. VT2e and c roughly express exercised capacity and capacity views, respectively, while acknowledging intentionality as a uniting feature necessary for grounding responsibility ascriptions. VT2e and c draw the lines of the debate more clearly while largely carving up the positions in the debate in the way they have previously been understood. While one may satisfy VT2c while *exercising* little, if any, voluntary control, for instance, the same is true of VT1c, so this is not a mark against VT2c. Similarly, one's omission of an intentional action is typically *susceptible* to one's voluntary control, as VT1c requires. Additionally, the shift in discussion to intentional actions is not unprecedented, as some have understood voluntary control in terms of intentional actions (Alston 1988; Hieronymi 2004), and talk of the voluntary can often be helpfully rephrased in terms of intentionality. For instance, Holly Smith writes, "the voluntary is that which *proceeds from, and carries out, the will*" (2011: 125). This might be rewritten: the voluntary is that which proceeds from some intention and carries out some action plan sustained or guided by the intention.

Beyond this precedent, however, there is good reason to think VT2e and c better capture what volitionists have really been discussing anyway. Though volitionists and nonvolitionists have typically focused on overt actions and omissions or attitudes,

mental actions, such as decisions, must also be accounted for. Decisions are intentional actions, but they are less clearly within one's voluntary control. Yet volitionists surely want to include decisions within the scope of items for which one can be directly responsible. My theses provide a clear way to do this in the manner in which volitionists are already inclined to do, and they provide clear justification for doing so. Intentional action is broader than voluntary control, but volitionists seem to want this broader focus anyway. The lack of attention to more challenging items like decisions has obscured this fact. I have simply widened the scope of the discussion because volitionists are interested in broader discussion and because the narrower scope has not led to productive results.

In sum, then, VT2e and c do not amount to an entire shift in topic such that they are no longer relevant to the volitionist debate. Voluntary control and intentional actions and omissions are closely related, so while there is some shift in focus, the shift is useful and precedented. The shift clarifies the differences between volitionists and nonvolitionists while preserving the way the conceptual territory has thus far been carved. It better captures certain items, including mental actions like decisions, as within the scope of direct responsibility, just as volitionists want but have not been able to clearly do thus far. Overall, my proposal does not stray far from the initial debate and yet it makes no appeal to fuzzy terms like "voluntary control" or mysterious notions of "the will".[28]

Perhaps one will think, then, that clarity is a virtue, but clarity isn't enough. How is my proposal any more illuminating or useful in advancing the debate than VT1e and c?

First, clarity itself can advance the debate. VT2e and c can point to new areas for exploration that will push the volitionist debate forward in ways that the traditional focus has not done and perhaps cannot do. My proposal ensures that those in the debate (and even those new to it) understand its parameters. For instance, Fischer and Tognazzini dismiss Angela Smith's nonvolitionist critique of tracing cases by claiming that the volitionist views she attacks are too narrow in their focus. Such views place emphasis on choice rather than control:

> Whereas a volitional approach requires that we trace back (directly or indirectly) to a *choice* in order to justify attributions of moral responsibility, a control model requires that we [trace] back (directly or indirectly) to an exercise of control, where such an exercise of control might be a choice, and [*sic*] action, or an omission. (2009: 549)

It's understandable why Fischer and Tognazzini understand volitionism this way, given Smith's emphasis on choice and inaptly titled "prior choice" views. Yet because she also appeals to voluntary control, a charitable interpretation understands "choice" more broadly, as an intentional undertaking. The obscurity surrounding voluntary control and choice can suggest that the debate is narrower than it is.

VT2e and c better illustrate that the debate concerns "choice" in a broad sense, not just as the upshot of deliberation. They also avoid misunderstandings with different conceptions of "voluntary control". For while Fischer and Tognazzini take a narrow view, Michael Robinson suggests a broad understanding of voluntary control in

[28] Thanks to an anonymous referee for pressing me to say more to address this concern.

discussing Frankfurt cases. Robinson criticizes McKenna and Fischer for writing as if "actions (particularly certain mental actions, such as deciding) are the only things over which agents have direct voluntary control" (2014: 439). We can have direct voluntary control over omissions as well, Robinson claims, such as when he simply omits to decide to get on a bus rather than choosing to get on it or choosing to not get on it (439–440). (Robinson seems to advocate VT2c.) VT2e and c can ensure everyone is having the same debate.

Besides this benefit, VT2e and c also help illustrate that the heart of the dispute is about the nature of obligations. Volitionists insist that it isn't fair to directly blame agents for, say, having or lacking some attitude, because this isn't under one's direct voluntary control. Nonvolitionists claim such blame is fair because the presence or absence of such attitudes express agential commitments. The debate typically devolves into discussions of types of responsibility, with nonvolitionists advocating answerability and volitionists insisting on accountability. Little progress is made, and it seems as if there is simply a verbal dispute or a dispute over the nature of responsibility.

I don't think either is quite right. Recall that according to Watson, accountability concerns demands and obligations. A tempting thought for volitionists is that the reactive attitudes are not justified in response to certain objectionable attitudes, noticing, neglecting, or forgetting, because these items are not under one's direct voluntary control and so one cannot be obligated to have some attitude, notice something, or remember something.

Why can one not be obligated to have some attitude or remember something? One might think that such an obligation would involve the denial of the intuitive principle that "ought" implies "can":

OIC: For any agent, *S*, if *S* ought to *A*, then *S* can *A*.

I cannot be obligated to have some attitude, the thought goes, because having such an attitude is not under my direct voluntary control. I cannot simply have some attitude; I can only try to bring the attitude about. Since nonvolitionists think it can be justified to directly blame someone for having an attitude, even though having the attitude isn't under one's direct voluntary control, it seems nonvolitionists must deny the intuitive OIC. Consequently, volitionists might think that nonvolitionists *cannot* understand responsibility as accountability.

There are various senses of "can", however, and it's not clear that nonvolitionists *must* deny OIC. If the "can" in OIC denotes *rational* ability rather than *voluntary* ability, it is false that I cannot have some particular attitude. For instance, one is *capable* of believing *p* rationally, even if one fails to do so. And here is where the debate seems to turn: what sorts of capacities are required for direct moral responsibility? Do agents need *voluntary control*, or is *rational control* sufficient? Volitionists protest that rational control is not sufficiently robust, and nonvolitionists disagree. The debate stalls.

Perhaps my proposal can get us going again. Part of the problem is that the scope of OIC is unclear. Consider this more detailed formulation:

OIC:* For any agent, *S*, and action, *A*, if *S* ought to do *A*, then *S* can do *A* (adapted from Henne et al. 2016: 283).

OIC* reveals that only actions may be substituted for *A*. Though often neglected, we can plausibly expand this to include omissions. But now we can see that OIC and OIC* don't create problems, but a closely related principle does:

OIA: For any agent, *S*, if *S* ought to *A*, then *A* is an action or omission.

There is good reason to think that the nature and scope of obligation itself, as expressed in OIA, are the heart of the issue rather than some version of OIC. Purported counterexamples that are typically offered to OIC involve actions or omissions.[29] For example, the due date for a class paper is Friday. Nevertheless, a student in the course procrastinates so that Friday morning, she cannot complete the paper on time. The student ought to turn in the paper, even though she cannot (Henderson 1966: 105). Such counterexamples focus on whether an obligation to act or omit can persist once an ability is lost. It seems those debating OIC already assume that obligations apply only to actions or omissions, but disagree over which actions and omissions one can be obligated to perform.

If this is right, nonvolitionists are not simply denying OIC (they may or may not); they are claiming that obligations can apply to more than just actions and omissions. Even if OIC is false, it may be that obligations apply only to actions and omissions. Nonvolitionists, then, may agree with volitionists that the reactive attitudes apply only in cases where obligations have been violated, but understand obligation differently than volitionists. According to the nonvolitionist, agents can be obligated to have certain attitudes, be certain ways, or remember certain things, and these obligations do not reduce to obligations to act or omit. Smith reveals this commitment in her response to Levy 2005:

> [H]e acknowledges that the thing I might reasonably have been expected to do cannot be 'to have remembered my friend's birthday,' for that is precisely the claim that is in dispute here. But it seems to me that this is precisely what the demand comes to in this case: As a friend, it is reasonable to expect me to remember my friends' birthdays (and not simply to do things to make it more probable that I will remember my friends' birthdays). (2008: 383, n. 19)

We might say that the volitionist thinks that our basic moral demand is to be *shown* good will from others, while the nonvolitionist thinks that our basic moral demand is for others to *have* good will for us.[30] Of course, nonvolitionists might also think that there is a demand that others show us good will, but the *fundamental* demand is simply to *have* this good will.

Importantly, then, what justifies any token reactive attitude for the nonvolitionist will be different from that which justifies the attitude for the volitionist. This difference is rooted in the different way in which they understand obligations, which becomes clearer after reframing volitionism in terms of intentional actions and omissions. When we think of volitionism in terms of voluntary control, the debate tends to converge on whether such control is needed for moral responsibility or not, and whether the reactive

[29] See, for instance, Henderson 1966, Zimmerman 1996, and Henne et al. 2016.

[30] Thanks to Robert Wallace for this insight and for useful discussion.

attitudes are fair or justified in response, and whether responsibility is accountability or something else. While these questions are absolutely worth exploring, the volitionist debate might see new life by shifting the framework of the question from voluntary control to intentional actions and omissions.[31]

Importantly, we cannot settle the debate by appealing to kinds of responsibility, and the focus on obligation will allow us to see why. As I noted in 1.2, many see the debate as a merely verbal dispute because each party is operating with a different notion of responsibility. But at least some nonvolitionists, like Angela Smith, reject these distinctions in responsibility and take themselves to be concerned with the same sort of responsibility as volitionists. Smith explicitly rejects attributability as a kind of responsibility (2008: 375–380), and she similarly rejects a distinction between answerability and accountability (2012: 586–589). The reason why she can do this, I suggest, is because she sees obligations as applying not only to intentional actions and omissions, but to attitudes as well.

Those who embrace different kinds of responsibility typically explain these differences in kind in virtue of the sorts of offenses committed and the sorts of responses warranted by such offenses. The sort of response warranted depends on the sort of offense committed. Watson suggests this when he writes, "Holding people responsible…involves a social setting in which we demand (require) certain conduct from one another and respond adversely to one another's failures to comply with these demands" (2004: 262). And in discussing accountability, Watson writes, "To be 'on the hook' in these and other cases is to be liable to certain reactions as a result of failing to do what one is required. To require or demand certain behavior of an agent is to lay it down that unless the agent so behaves she will be liable to certain adverse or unwelcome treatment" (275). These passages reveal that as Watson understands accountability, it involves being subject to the reactive attitudes in virtue of failing to meet certain obligations.[32]

More recently, David Shoemaker has defended a tripartite distinction in kinds of responsibility, embracing attributability, answerability, and accountability (2011). Yet as Smith understands Shoemaker, "answerability appears to differ from accountability primarily on the grounds that answerability licenses only aretaic appraisals and other nonsanctioning modifications to one's actions and attitudes toward the agent, while accountability licenses more explicit forms of sanctioning activity" (Smith 2012: 586–587). According to Shoemaker, the reason for this difference lies in the sorts of

[31] To clarify, I am not suggesting that understanding the volitionist theses in terms of intentional actions and omissions is the only way to achieve this upshot or that we must understand volitionism in this way to see that obligation is a useful path to explore in advancing the debate. One purpose for understanding volitionism in terms of intentional actions and omissions is clarity, but another contingent benefit is that the disagreement regarding the scope of obligation becomes clearer. This is in part because we begin to think about the debate in a different, yet related way. Thus, even if one rejects VT2e and c, one may nevertheless agree that the nature and scope of obligation is the key point of disagreement between volitionists and nonvolitionists. Conversely, one may disagree that turning to obligation is a worthwhile avenue forward, yet agree that VT2e and c sharpen and clarify the boundaries of the volitionist debate. Here, I am arguing both for a different understanding of the volitionist theses and for a new way forward in the debate. While these two projects are not linked by necessity, the former nudges us toward the latter, and having the clearer distinctions in mind as we discuss the nature and scope of obligation will provide a more promising discussion. I am grateful to an anonymous referee for urging me to say more on this point.

[32] Notice that Watson discusses only conduct or behavior; the idea that one could be obligated simply to have certain attitudes or to be certain ways is absent from this understanding of accountability.

normative demands of various relationships. Shoemaker distinguishes between two cases to make his point—the former involving answerability and the latter involving accountability. In *Anniversary*, year after year George fails to pick up on his wife Martha's hints of what she wants for her anniversary, instead buying her carnations (her least favorite flower). After several years, Martha stops dropping hints and resigns herself to disappointment. In *Cheating*, George is unfaithful to Martha and engages in several extramarital affairs, and Martha angrily takes a swing at him with a golf club (2011: 620–621). Only in *Cheating,* says Shoemaker, is Martha's resentment fitting:

> Martha has a right to George's fidelity, given the mutual claims the parties to the relationship have made on one another in the very formation of the relationship (perhaps explicitly, in the case of shared marriage vows); Martha does not, however, have a right to George's sensitivity to her hints with respect to gift-giving…. So in letting her down in Anniversary, he has not done anything to render resentment fitting, for he did not violate any of her claims on him; rather, he simply dashed her hopes (622).

Again, notice that according to Shoemaker, resentment is fitting because a moral demand has been flouted, whereas Martha has no legitimate demand that George be sensitive to her hints. The sort of response justified is a function of whether one has violated some obligation, and, one might argue, an individual has no obligation to be sensitive to hints in this way.

Smith, of course, rejects that the different responses in *Anniversary* and *Cheating* are a function of different kinds of responsibility. Instead, she argues, the difference in the seriousness of the failures warrants a different response, but different responses do not indicate different kinds of responsibility (2012: 587).[33] One could just as easily imagine Martha expressing her resentment in *Anniversary* by angrily exclaiming, "Why do you bother, George?—don't you see, I simply don't like carnations!" (588). And, according to Smith, Martha's resentment seems justified. How can this be? Smith does not elaborate, but one plausible explanation is that Smith *does* see George as failing in some obligation to notice what his wife wants, just as Smith thinks the demand in her birthday forgetting case is simply to have remembered her friend's birthday.

The sort of responsibility at issue, then, depends on how we understand obligation. The different ways in which volitionists and nonvolitionists understand the nature and scope of obligation informs their understanding of moral responsibility and the nature of the reactive attitudes. Certain responses, such as the reactive attitudes, are warranted only when (and because) one has violated some obligation. Thus, we cannot simply dismiss the debate as a case of verbal dispute until we discuss the nature and scope of

[33] It is important that we keep in mind that the moral responses volitionists and nonvolitionists discuss are "key elements of our actual moral practices—aretaic appraisals; feelings of disappointment, resentment, or indignation; modifications for our attitudes and actions in response to perceived relational impairments; explicit acts of reproach or censure" (Smith 2012: 588). If we up the ante significantly, Smith says, such that we understand "'being morally responsible for X' as 'being a legitimate target of eternal damnation or reward on the basis of X'" our understanding of responsibility might change (588). So the responses warranted might, at some point, affect our understanding of moral responsibility, but in this discussion, they will not. Such a strong response is not at issue here, and it is not clear that many theorists working on free will or moral responsibility really mean to rely on responsibility that warrants eternal punishment (Clarke 2005: 20–22). Perhaps this is because no sort of obligation violation could merit such responses.

obligation. Volitionists insist that one is not obligated to simply have (or lack) some attitude, so the reactive attitudes are not a justified response to the having of some attitude. Nonvolitionists see obligation as applying even to attitudes. Volitionists may understand responsibility as accountability in virtue of thinking that the negative reactive attitudes require some obligation violation to be justified. Nonvolitionists understand responsibility as answerability because they can allow that obligations apply to attitudes as well. The question, then, is not the sorts of responses that are justified, but the nature and scope of obligation itself, the violation or surpassing of which might merit some reactive attitude.

One might think that I have the relationship between responsibility and obligation backwards. I have suggested that one's understanding of obligation informs one's understanding of responsibility, but perhaps one's understanding of responsibility informs one's understanding of obligation instead.[34] I have been suggesting that obligation is more fundamental than responsibility, but perhaps responsibility is more fundamental than obligation.

Perhaps this relationship goes both ways, such that one's understanding of obligation informs one's understanding of responsibility *and* one's understanding of responsibility informs one's understanding of obligation. But obligation seems to me to have the priority. As we have seen, accountability is typically understood and described in terms of obligations. Those who embrace different kinds of moral responsibility think of accountability as pertaining to obligation violations, unlike attributability. Given that responsibility is defined in terms of obligation, I take obligation to be the more fundamental concept.

The distinctions between kinds of responsibility are not always clear, which is why I prefer to focus on the reactive attitudes. A volitionist says that *S* deserves the reactive attitudes only when *S* violates an obligation. But *S* cannot be obligated to have certain attitudes, so *S* does not deserve blame for having that attitude. The nonvolitionist says that *S* can be obligated to have certain attitudes, so *S* may well deserve blame for having that attitude. Given that both volitionists and nonvolitionists are focused on responsibility that merits the reactive attitudes, this must be a robust moral 'ought'—not just an 'ought' of goodness, for instance. The 'ought' of goodness does not invoke or justify the reactive attitudes. So while one's understanding of responsibility and obligation may be interdependent in certain ways, given the concern with the reactive attitudes, I take it that both volitionists and nonvolitionists are concerned with the same robust moral 'ought'. The difference for volitionists and nonvolitionists is in what agents can be obligated, in this robust sense, to do, be, or have. If one's understanding of responsibility grounded one's understanding of obligation and the type of 'ought' at issue, I do not see how we could even make sense of different kinds of responsibility, and I would expect to see volitionists and nonvolitionists using different types of moral 'ought's.

If I am right that the understanding of obligation informs the understanding of responsibility and that volitionists and nonvolitionists are largely disputing what obligations agents have, this provides a way forward in the debate: we must explore the nature and scope of obligation itself. Rather than argue about the true nature of moral responsibility we might explore whether blaming for items other than intentional

[34] Thanks to an anonymous referee for raising this worry.

actions and omissions even *makes sense*. Thus, the challenge for nonvolitionists is showing how and why obligations can apply to things other than actions and omissions. If obligations can apply to more than just actions and omissions, nonvolitionism may be a viable contender as a theory of responsibility. If not, nonvolitionists have cause to be concerned.

Perhaps this avenue will ultimately bear no fruit or lead us back to a discussion of voluntary control. But it is not clear that this is what must happen, because this is not a debate we have had. The only way to know is to have the debate. Given the lack of progress in the volitionist debate thus far, it is worth asking whether we are even having the debate in the right way. My suggestion is shifting our focus from voluntary control and kinds of responsibility to the scope of responsibility, grounded out in terms of obligations by way of intentional actions. If we shift our focus in this way and make progress, that will be a positive result. But even if we shift our focus and discover this avenue is a dead end, perhaps the reasoning behind why this avenue is fruitless will reveal something more significant. At the very least, we can be more confident that the debate cannot be resolved without tackling the challenging questions of the nature of voluntary control.

It is worth summarizing the discussion. I am suggesting a different approach in the debate between volitionists and nonvolitionists. First, we must be clearer about what the debate itself is. Most have understood the debate to be about the scope of responsibility, somehow related to control, but the concepts and lines are fuzzy. To clarify the debate, I have suggested that volitionists think that we can be directly responsible only for intentional actions and omissions, while nonvolitionists think that we can be directly responsible for intentional actions and omissions as well as attitudes. Second, and relatedly, we must be clearer about what questions must be answered to progress in the debate. Thus far, the debate has centered around the idea that the disagreement is due to the control requirement of responsibility. But instead of shifting from the scope of what we are responsible for to the control requirement and the nature of this requirement, we could shift from the scope of responsibility to the scope and nature of obligation. Hopefully this shift can provide the progress the previous way of understanding the debate could not.

6 Conclusion

To close, reconsider *Airport*. Volitionists who favor VT2e claim that Josh is not directly responsible for forgetting to pick up Kate. His forgetting is neither an intentional action nor an intentional omission. Josh is not obligated to remember something, but rather to bring it about that he remember something. Consequently, Josh may be derivatively responsible for forgetting to pick up Kate, but not directly responsible for this. For similar reasons, Josh may be derivatively responsible for failing to pick up Kate, but not directly responsible for so failing.

Interestingly, a proponent of VT2c could claim that Josh is directly responsible for failing to pick up Kate, as this is an omission of an intentional action. Because Josh promised to pick up Kate, he is obligated to do so. Yet Josh is not directly responsible for forgetting to pick up Kate, as forgetting is neither an intentional action nor an omission of an intentional action. It is not something to which obligations can rightfully

apply, according to the volitionist. But he may nevertheless be directly responsible for failing to pick up Kate.

Nonvolitionists may hold that Josh is directly responsible for both forgetting to pick up Kate and for failing to pick up Kate, because they do not restrict direct responsibility to intentional actions or omissions. Josh is obligated to remember to pick up Kate and to pick her up, one might think, in virtue of his relationship with her and his promise. In failing to remember and in failing to pick up Kate, a nonvolitionist might argue, Josh expresses an objectionable attitude, or perhaps reflects that he lacks attitudes he ought to have regarding Kate.

Neither volitionists nor nonvolitionists have carefully articulated a volitionist thesis or explained voluntary control. Such tasks are deceptively difficult. I have suggested we shift our focus to VT2e and c instead. While I haven't argued for or against the truth of either thesis, I have argued that they make salient that the heart of the disagreement between volitionists and nonvolitionists regards the nature of obligations themselves. This understanding clarifies the dispute between volitionists and nonvolitionists and helps us understand not only why there is disagreement, but also how to make progress toward resolving it.

Acknowledgments Thanks to Randolph Clarke, Alfred Mele, and Daniel Miller for helpful feedback on previous drafts of this paper. For productive conversations, thanks also to John Schwenkler, Robert Wallace, and those in attendance for a presentation of a portion of this paper at the 2016 Pacific Division Meeting of the American Philosophical Association. Finally, thanks to an anonymous referee for thoughtful questions and objections on a previous draft of the paper.

References

Adams, R. (1985). Involuntary sins. *The Philosophical Review, 93*, 3–31.

Adams, F., & Mele, A. (1992). The intention/volition debate. *Canadian Journal of Philosophy, 22*, 323–337.

Alston, W. (1988). The deontological conception of epistemic justification. *Philosophical Perspectives, 2*, 257–299.

Chuard, P., & Southwood, N. (2009). Epistemic norms without voluntary control. *Nous, 43*, 599–632.

Clarke, R. (2005). On an argument for the impossibility of moral responsibility. *Midwest Studies in Philosophy, 29*, 13–24.

Clarke, R. (2014). *Omissions: Agency, metaphysics, and responsibility*. New York: Oxford University Press.

Clarke, R., & Reed, T. (2015). Free will and agential powers. In D. Shoemaker (Ed.), *Oxford studies in agency and responsibility, volume 3*. Oxford: Oxford University Press.

Clarke, R., McKenna, M., & Smith, A. (2015). Introduction. In R. Clark, M. McKenna, & A. Smith (Eds.), *The nature of moral responsibility: New essays*. New York: Oxford University Press.

Coates, D. J., & Tognazzini, N. (2012). The nature and ethics of blame. *Philosophy Compass, 7*(3), 197–207.

Fischer, J. M., & Tognazzini, N. (2009). The truth about tracing. *Nous, 43*, 531–556.

Ginet, C. (1990). *On action*. Cambridge: Cambridge University Press.

Henderson, G. (1966). 'Ought' implies 'can'. *Philosophy, 41*, 101–112.

Henne, P., Chituc, V., de Brigard, F., & Sinnott-Armstrong, W. (2016). An empirical refutation of 'ought' implies 'can'. *Analysis, 76*, 283–290.

Hieronymi, P. (2004). The force and fairness of blame. *Philosophical Perspectives, 18*, 115–148.

Hieronymi, P. (2008). Responsibility for believing. *Synthese, 161*, 357–373.

van Inwagen, P. (1983). *An essay on free will*. Oxford: Oxford University Press.

Kane, R. (1996). *The significance of free will*. Oxford: Oxford University Press.

Levy, N. (2005). The good, the bad, and the blameworthy. *Journal of Ethics and Social Philosophy, 1*, 2–16.

Levy, N. (2008). Restoring control: Comments on George Sher. *Philosophia, 36*, 213–221.

Levy, N. (2014). *Consciousness and moral responsibility*. Oxford: Oxford University Press.
Levy, N., & McKenna, M. (2009). Recent work on free will and moral responsibility. *Philosophy Compass, 4*(1), 96–133.
McCann, H. (1998). *The works of agency: On human action, will, and freedom*. Ithica: Cornell University Press.
McKenna, M. (2008). Putting the lie on the control condition for moral responsibility. *Philosophical Studies, 139*, 29–37.
McKenna, M. (2012). *Conversation and responsibility*. New York: Oxford University Press.
Mele, A. (2003). *Motivation and agency*. Oxford: Oxford University Press.
Mele, A. (2017). Direct control. *Philosophical Studies, 174*, 275–290.
Mele, A., & Moser, P. (1994). Intentional action. *Nous, 28*, 39–68.
O'Shaughnessy, B. (2008). *The will: A dual aspect theory*, 2nd ed., in 2 volumes. Cambridge: Cambridge University Press.
O'Shaughnessy, B. (2009). Trying and acting. In L. O'Brien and M. Soteriou (Eds.), *Mental actions* (pp. 163–172). Oxford: Oxford University Press.
Robinson, M. (2014). The limits of limited-blockage Frankfurt-style cases. *Philosophical Studies, 169*, 429–446.
Shepherd, J. (2015). Deciding as intentional action: Control over decisions. *Australasian Journal of Philosophy, 93*, 335–351.
Sher, G. (2009). *Who knew? Responsibility without awareness*. Oxford: Oxford University Press.
Shoemaker, D. (2011). Attributability, answerability, and accountability: Toward a wider theory of moral responsibility. *Ethics, 121*, 603–632.
Smith, A. (2005). Responsibility for attitudes: Activity and passivity in mental life. *Ethics, 115*, 236–271.
Smith, A. (2007). On being responsible and holding responsible. *The Journal of Ethics, 11*(4), 465–484.
Smith, A. (2008). Control, responsibility, and moral assessment. *Philosophical Studies, 138*, 367–392.
Smith, H. (2011). Non-tracing cases of culpable ignorance. *Criminal Law and Philosophy, 5*, 115–146.
Smith, A. (2012). Attributability, answerability, and accountability: In defense of a unified account. *Ethics, 122*, 575–589.
Vargas, M. (2005). The trouble with tracing. *Midwest Studies in Philosophy, 29*, 269–291.
Watson, G. (2004). Two faces of responsibility. In *Agency and answerability: Selected essays* (pp. 260–288). Oxford: Clarendon Press.
Zimmerman, M. (1988). *An essay on moral responsibility*. Totowa, NJ: Rowman and Littlefield.
Zimmerman, M. (1996). *The concept of moral obligation*. New York: Cambridge University Press.