

# In defense of doxastic blame

Lindsay Rettler<sup>1</sup>

Received: 8 February 2016 / Accepted: 4 February 2017  
© Springer Science+Business Media Dordrecht 2017

**Abstract** In this paper I articulate a view of doxastic control that helps defend the legitimacy of our practice of blaming people for their beliefs. I distinguish between three types of doxastic control: intention-based, reason-based, and influence-based. First I argue that, although we lack direct intention-based control over our beliefs, such control is not necessary for legitimate doxastic blame. Second, I suggest that we distinguish two types of reason-responsiveness: sensitivity to reasons and appreciation of reasons. I argue that while both capacities are necessary for satisfying the control condition, neither is sufficient. Finally, I defend an influence-based view of doxastic control according to which we have the capacity to execute intentions to engage in reflection that causally influences our beliefs in positive epistemic ways. This capacity is both necessary and sufficient for satisfying the control condition for legitimate doxastic blame. I end by defending the view from two objections: that reflection is not necessary for meeting the control condition, and that it is not sufficient.

**Keywords** Epistemology · Doxastic agency · Doxastic blame · Ethics of belief · Doxastic voluntarism · Moral blame

## Introduction

We often hold people responsible for their beliefs. We blame people for believing what they shouldn't, just as we blame people for *doing* what they shouldn't. However, many philosophers have argued that we don't have the same control over our beliefs that we

---

✉ Lindsay Rettler  
lindsayrettler@gmail.com

<sup>1</sup> Baylor University, Waco, TX, USA

have over our actions.<sup>1</sup> For example, while we can turn on the television at will, we cannot believe at will that the television is already on.<sup>2</sup> But if we lack control over our beliefs, then it's difficult to see how we could be legitimately blamed for them.

The challenge for doxastic blame can be formalized as the *Anti-Blame Argument*:

- (1) If agents are legitimately subject to blame for their beliefs, then they have control over their beliefs.
- (2) It is *not* the case that agents have control over their beliefs.
- (3) Therefore, agents are not legitimately subject to blame for their beliefs.

There are two ways to resist the Anti-Blame Argument. The first is to deny premise 1 and argue that legitimate doxastic blame does *not* require control. The second is to agree that legitimate doxastic blame requires control, and to argue, contra premise 2, that we *do* have the requisite control. Of course, there is a *third* option available: one could accept the argument as sound and agree that doxastic blame is not legitimate. But given that our practices of doxastic blame are deeply entrenched, this would be quite revisionary.<sup>3</sup>

Which option one takes depends on how one construes both the notion of control *and* the notion of blame. Richard Feldman defends the first option—that legitimate doxastic blame doesn't require control.<sup>4</sup> However, as I will argue, this option is not viable once we clarify the nature of doxastic blame.

Some other philosophers have opted for the second way of defending the legitimacy of doxastic blame: they claim that we do have the requisite kind of control over our beliefs. In particular, they appeal to the notion of reason-responsiveness to satisfy the control condition for legitimate doxastic blame. While I agree that we should reject premise 2 of the Anti-Blame Argument, I will argue that given the nature of doxastic blame, reason-responsiveness itself is not sufficient to satisfy the control condition for legitimate doxastic blame. In turn, I will advocate an alternative account of doxastic control that is sufficient. While I defend the claim that we have control over our beliefs, this does not make me a doxastic voluntarist given standard definitions of voluntarism, according to which we have direct intention-based control over our beliefs.<sup>5</sup>

I begin, in Sect. 1, with a few preliminaries. Then in Sect. 2 I explain in detail what doxastic blame is, and why I dismiss the view that legitimate doxastic blame does not require control. Following this, I evaluate three types of doxastic control: *intention-based*, *reason-based*, and *reflective control*. In Sect. 3 I argue that though we lack intention-based control over our beliefs, such control is *not necessary* for legitimate

<sup>1</sup> For example, Williams (1973, p. 148) famously argues that we cannot believe at will. Shah (2002, pp. 440–441) and Weatherson (2008, pp. 554–555) argue that we cannot believe for practical reasons. Hieronymi (2006, 2009) argues we cannot believe for what she calls extrinsic reasons.

<sup>2</sup> The difference can be put in terms of reasons: while we can act for any reason that counts in favor of acting, it is not the case that we can believe for any reason that counts in favor of believing. See, for example, Hieronymi (2006, pp. 45–46).

<sup>3</sup> See Woudenberg (2009) for a good discussion of just how revisionary it would be to conclude that doxastic blame is illegitimate.

<sup>4</sup> See Feldman (2000, 2001). I discuss his view further in Sect. 2.2.

<sup>5</sup> I argue against the view that we have such intention-based doxastic control in Sect. 3.

doxastic blame. In Sect. 4 I argue that though we do have reason-based control, such control is *not sufficient*. In Sect. 5 I present my preferred view of doxastic control, the *Reflective Control View*, which advocates a form of influence-based control. I argue that indirect reflective control is both necessary and sufficient to satisfy the control condition for legitimate doxastic blame. I end, in Sect. 6, by defending the view from two objections.

## 1 Preliminaries

I start with three clarifications of the scope of my project: First, for any agent *S*, if *S* is legitimately blamed for *x*, then *S* has, or at some time had, control over *x*. There may be some attenuated notions of blame that do not require control. For example, blaming the Texas weather for frying my laptop contrasts with blaming my friend for carelessly dumping her coffee on it. The former is, at best, a weak sense of blame; it is the latter robust sense of blame that I am concerned with, and I will say more about it in the following section.

Second, control is required for legitimate blame, but it's not the only condition that's necessary. Other plausible requirements for legitimate blame include an *attribution condition*, a *value condition*, and an *epistemic condition*. For example, if an agent, *S*, is legitimately blamed for *x*, then *x* can be properly attributed to *S*, *x* is wrong in some sense, and *S* knew or could have known that *x* is wrong.<sup>6</sup> Here, I focus exclusively on the capacity that is both necessary and sufficient for fulfilling the *control condition* for legitimate doxastic blame.

The third clarification concerns how my project relates to defending the deontological conception of epistemic justification, according to which epistemic justification is analyzed in deontological concepts—concepts such as obligation, prohibition, and blame. William Alston famously argues that because we lack voluntary control over our beliefs, we *cannot* analyze epistemic justification in any of these terms.<sup>7</sup> While I will argue that we do have a form of control over what we believe, it is not my goal to vindicate the deontological conception of justification. According to the version of doxastic control that I will defend, we have control over only *some* of our beliefs. But since *all* beliefs can be evaluated with respect to whether they are epistemically justified, I agree with Alston that we should reject the deontological conception of epistemic justification.<sup>8</sup>

## 2 Doxastic blame

I now turn to the nature of doxastic blame. Consider three examples:

<sup>6</sup> See also Wolf's discussion of what she calls the "condition of freedom" and the "condition of value," both of which must be met in order for a person to be morally responsible for her action (1980, p. 151).

<sup>7</sup> Alston (1988); see also (1985).

<sup>8</sup> Note that if Feldman (2000, 2001) is right when he claims that agents can still be subject to role obligations even if they lack control over their beliefs, then conceiving of epistemic justification in terms of role obligations is a live option. The stronger versions of deontological justification—conceived of in terms of responsibility and blame—are the ones that I reject, along with Alston.

CASE A: Aaron blames Sam for believing that he, Aaron, ate the last of Sam's cereal. Aaron resents Sam for believing that he is the culprit, given that they live with two other roommates whom Sam has equal reason to believe guilty of the crime. Aaron resolves to make himself scarce whenever Sam is around.

CASE B: As a philosophy professor who regularly teaches critical thinking and logic courses, Brette finds herself indignant at the unreasonable beliefs of her incoming freshman class this fall. In particular, Brette blames her student Seth for believing that his chances of meeting a celebrity are high just because he's never met one until now.

CASE C: Cassia dislikes one of her interns and finds herself judging Danielle's work more harshly than she judges the work of the other interns. She concludes that Danielle is incompetent at her job, but when she's describing her reasons to her friend, her friend points out how prejudiced she's been in her interpretation of her evidence. Cassia then blames herself: she feels guilty for believing Danielle is incompetent, when all along she had good reason to judge her a fine worker.

In each of these cases the agent in question blames another person, or herself, for a belief that is faulty in some way. Notice that the blame in question is not a mere causal judgment that *this* person rather than someone else authored the belief. Nor does the blame target any harmful actions caused by the faulty beliefs.<sup>9</sup> Nor does the blame target what might be called *epistemic* actions, such as gathering the wrong evidence. Instead, the blame specifically targets the faulty belief of another agent. The blaming agents hold others responsible for their beliefs by feeling resentment, indignation, and guilt, respectively.

## 2.1 The nature of doxastic blame

To better understand the nature of this blame, consider for a moment the nature of moral blame for action. Many contemporary accounts of moral blame share in common the idea that moral blame is distinct from both mere *moral grading*—i.e. judging or assessing an agent relative to various moral standards of evaluation—and *sanctioning*—i.e. subjecting an agent to some form of harsh treatment in response to her failure. Moral blame is thought to have a distinctive characteristic force that mere negative evaluative assessment does not have, which indicates that grading is not sufficient for moral blame. As Pamela Hieronymi puts it:

Blame, it is thought, goes beyond simple description or mere grading...Being morally blamed involves a more serious sort of criticism than being told your vocal performance was flat, your cooking bland, your conversation dull, or your

<sup>9</sup> Consider Clifford's (1999, pp. 70–71) infamous case of doxastic blame in which a shipowner who falsely believes his ship is safe sends it out to sea, resulting in many deaths. Some have argued that this does not show that we blame people for their beliefs because the case is better understood in terms of blaming the shipowner for the actions caused by his belief. However, while we often do blame people for the actions resulting from their beliefs, we *also* sometimes blame them for their beliefs, as the cases above illustrate.

sentences opaque. Blame, unlike mere description, carries a characteristic depth, force, or sting.<sup>10</sup>

While characterizing blame in terms of sanctioning might seem at first glance to account for the force of blame—after all, sanctions are forceful—sanctioning is also insufficient for moral blame. One can sanction one’s pet without blaming it, as when one wishes to teach the pet to associate negative consequences with a particular behavior.

The force of blame can be explained by appealing to the fact that in blaming someone, we place a normative demand on her. The nature of the normative demand involved in moral blame becomes clearer when we note that moral blame functions to effect a certain kind of response from the one blamed. When I blame you for wantonly hurting people, I am not merely judging that you violated a moral norm to respect people—i.e. I’m not simply morally grading you. Neither am I necessarily sanctioning you. Rather, I’m demanding a response from you: I demand that you acknowledge your failure to treat people well. Doing this requires you to recognize that you failed to act on the moral reasons you had for treating people well.

Many accounts of moral blame emphasize this notion of a demand.<sup>11</sup> I remain neutral with respect to which account best characterizes the normative demand of blame. One might hold with P. F. Strawson and R. Jay Wallace that blame necessarily involves the reactive attitudes, which are intimately linked with the demand that others treat us with reasonable regard.<sup>12</sup> Such attitudes include resentment, indignation, and guilt. Alternatively, one might hold with T.M. Scanlon and Angela Smith that blame need not involve any reactive attitudes but instead involves any attitudes or intentions made appropriate by an agent’s moral failure.<sup>13</sup> In that case, one’s demand might be expressed by the decision to stop entrusting any more secrets to the friend who betrayed a confidence, or as in Case A, by Aaron’s resolve to make himself scarce whenever Sam is around. When these actions function as blame, they are undertaken as a way of expressing “implicit demands for justification, demands that are made on the basis of moral standards we expect reasonable persons to accept” (Smith 2008, p. 369). One can think of the demand as a normative expectation that the blamed agent acknowledge her failure. The term ‘expectation’, though, can be misleading since there is a predictive sense of expectation, according to which one believes that the blamed agent *will* acknowledge her failure, and this belief is not necessary for blame. One can demand that an agent acknowledge her failure while believing that she will not in fact do it.

Additionally, note that one might hold that the demand inherent in blame is always communicated—i.e. since the concept of a demand is inherently communicative, and

<sup>10</sup> Hieronymi (2004, pp. 116–117). See also: Scanlon (1988) reprinted in Watson (2003, pp. 365, 370) and Scanlon (1998, p. 269), Smith (2008, p. 369) and (2013, p. 29), Wallace (1994, pp. 80–81), and Wolf (1990, pp. 40–41, 64).

<sup>11</sup> See, for example, Hieronymi (2004, p. 117), McKenna (1998), Strawson (1974, pp. 6–7, 14–16), and Wallace (1994, 2007, 2010, 2011).

<sup>12</sup> See Strawson (1974) and Wallace (1994, 2011). See also Wolf (2011).

<sup>13</sup> See Scanlon (2008) and Smith (2013). See also Sher (2006).

since blame involves demand, blame must always be communicated. Alternatively, one might hold that one can make a demand without communicating it, and therefore one can blame someone without communicating it to the person. Both are legitimate views, and I do not take a stand here on which is correct.

The fact that blame involves a normative demand opens up the possibility that there are varieties of blame in addition to moral blame. Indeed, I think this is indicated by our practices. While we often blame individuals for violating distinctively moral norms, we also blame individuals for violating other sorts of norms. These norms include the norms of wisdom, norms of etiquette, norms of efficiency, and so on. Sometimes we merely criticize agents for failing to act in accordance with such norms—we grade them. For example, I judge that you slighted me at a recent social function. But other times we additionally blame agents for violating such norms, as when I resent you for slighting me. In so resenting you, I demand that you acknowledge your failure of etiquette—I demand that you recognize that you failed to act on the reasons you had to be polite to me.

More important for my purposes here is that we sometimes blame people in this same sense for violating the norms that govern *doxastic attitudes*, such as belief, disbelief, and suspension of judgment. The norms that govern doxastic attitudes are *epistemic norms*, so we might call this blame *epistemic blame*. While I remain neutral on what makes a failure distinctively epistemic, examples include having doxastic attitudes that are unjustified, unsupported, unreasonable, inconsistent, incoherent, and so forth. As I am conceiving of it, epistemic blame is a species of doxastic blame; it targets doxastic attitudes that violate epistemic norms. If there are other forms of doxastic blame, they target doxastic attitudes that violate other kinds of norms, such as moral norms. I remain open to the possibility that there are distinctively moral norms of belief. However, for simplicity I focus here on doxastic blame that targets violations of epistemic norms. Also, for simplicity I focus on blaming people for beliefs in particular, though we can certainly blame people for a wide range of doxastic attitudes. What I say about belief applies equally to other doxastic attitudes.

In summary, epistemic blame is a response that amounts to holding an individual responsible for failing to meet an epistemic standard. This holding responsible essentially involves demanding that the blamed agent acknowledge her failure to believe correctly because she acknowledges the epistemic reasons she rejected or ignored in forming or maintaining her belief.

## 2.2 Doxastic blame and control

We're now in a position to see why the first response to the Anti-Blame Argument fails. As I mentioned earlier, one could try to resist the argument by maintaining that legitimate doxastic blame does not require control. Richard Feldman suggests that doxastic blame can be legitimate, independent of control, if we characterize our doxastic obligations as *role obligations*. According to Feldman, role obligations simply “describe the right way to play certain roles” (2000, p. 676). Just as a teacher ought to explain information clearly and a parent ought to take care of her children, so also a believer ought to believe in accordance with her evidence. According to Feldman,

“It is our plight to be believers. We ought to do it right. It doesn’t matter that in some cases we are unable to do so” (2000, p. 676).

While Feldman’s primary goal is to demonstrate that agents can legitimately be subject to doxastic obligations even if they lack control over whether they fulfill them, Feldman suggests that we can also explain legitimate doxastic blame with reference to these role obligations.<sup>14</sup> The idea is that regardless of control, an individual can be legitimately blamed for failing her doxastic obligations in virtue of the role she occupies. In support of this Feldman points out that “we do praise and blame people for attributes, such as beauty, that they are unable to control” (2000, p. 676).

I suggest, however, that when we heed the distinction between blame and merely evaluative assessment, i.e. grading, we see both that people cannot be legitimately blamed for physical attributes over which they have no control, and that people cannot be legitimately blamed for violating a role obligation over which they have no control.<sup>15</sup> In both of these cases, merely evaluative assessment, whether negative or positive, is certainly legitimate. This is because such assessment is legitimate independent of whether an agent has control. I can legitimately judge people as too short or as unattractive, regardless of whether they have control over their height or appearance. Such assessments merely compare or contrast an individual with some particular evaluative standard. But this is distinct from blaming the person for failing to meet various height or appearance standards.

If  $S$  is legitimately blamed for her height, then  $S$  has, or at some time had, control over her height; likewise with respect to the occupying of a role. Teaching and parenting are both roles that individuals typically *do* have some control over whether they occupy, and once they do, most people have some control over how well they fulfill the role. Indeed this is why it seems legitimate to blame teachers and parents for failure to parent or teach well. However, the role of believer is not something that people have any control over whether they occupy.<sup>16</sup> So if doxastic blame is legitimate when people fail to occupy the role well, then they must, at some time, have some control over how well they occupy it. In other words, for doxastic blame to be legitimate, individuals must have some kind of control over their beliefs.<sup>17</sup> We now turn to three candidates for such control.

### 3 Direct intention-based control

The first candidate is *direct intention-based control*. An agent has direct intention-based control over  $\phi$ -ing just in case she can  $\phi$  directly as the result of an intention to

<sup>14</sup> See Feldman (2000, p. 676) and Feldman (2001, pp. 89–90).

<sup>15</sup> Peels (2013, pp. 8–9) offers a similar argument against Feldman’s claim that agents can be subject to doxastic obligations even if they lack control.

<sup>16</sup> Of course, people do have control over whether they occupy the role of believer insofar as they have control over whether they are alive. But this is a form of control people have over almost everything, and as such is not very interesting.

<sup>17</sup> Note that it may well be that being subject to epistemic obligation does not require that one has control. But being legitimately blamed for violating an epistemic obligation (or for failing to realize various epistemic values) does require that an agent has some form of control.

$\phi$ . We have this sort of control over our actions. The doxastic parallel is the ability to believe a proposition,  $p$ , as a direct result of an intention to believe that  $p$ .

### 3.1 Do we have direct intention-based control?

William Alston provides the most well-known argument that we *cannot* control our beliefs via directly executing intentions to believe specific propositions.<sup>18</sup> To begin, Alston simply points out that we can't believe something as a result of an intention to believe it "just like that" (1988, p. 263).<sup>19</sup> If we try, for example, to immediately execute an intention to believe that the U.S. is still a colony of Great Britain, we don't succeed.<sup>20</sup> But it might seem plausible that though we can't directly intend to believe that  $p$  in many cases, we can do so in cases in which the evidence does not seem to clearly support one proposition, such as matters of politics or religion. In such cases, people speak of deciding or choosing what to believe. Even here, though, Alston thinks we do not form a specific belief as a direct result of an intention to do so. He argues that either the direct cause of the belief is how the evidence seems to us at that moment, or else we don't in fact form the belief that  $p$  but merely accept  $p$  or act as if  $p$  is true.<sup>21</sup>

I think Alston is right that we lack direct intention-based doxastic control over all of our beliefs. Note that Alston doesn't argue that it's impossible; rather, he defends the view that we contingently lack this ability. Others have defended the stronger claim that as a matter of conceptual necessity we lack the ability to directly execute an intention to believe that  $p$ .<sup>22</sup> In any case, whether it's contingent or necessary, I conclude that we lack direct intention-based control over our beliefs.

### 3.2 Is direct intention-based control necessary for doxastic control?

Given that we lack direct intention-based control, if it were required to satisfy the control condition for legitimate doxastic blame, then doxastic blame would indeed be illegitimate. However, while the ability to  $\phi$  directly as a result of an intention might be sufficient for control, it is not necessary.

Consider that we are legitimately blamed for a variety of things over which we lack direct intention-based control. Such things include various complex actions, omissions, emotions, desires, and character traits, including vices. For example, we can be legitimately blamed for arriving late to an appointment, even though arriving promptly is mediated by our ability to carry out an intention to leave for the appointment with enough travel time. We can be legitimately blamed for our poor health, even though

<sup>18</sup> See Alston (1988, pp. 263–277).

<sup>19</sup> Sometimes believing directly as a result of an intention is referred to as "believing at will" or "deciding to believe". However, there is no agreed upon definition of these locutions, so I think it can be misleading to use them without specifying the phenomenon one has in mind.

<sup>20</sup> Alston (1988, p. 263).

<sup>21</sup> Alston (1988, pp. 264–268).

<sup>22</sup> See Buckareff (2014, pp. 33–40), Hieronymi (2005, 2006, 2009), Kelly (2002), Setiya (2008) and Shah (2006).



the state of our health depends on our prior ability to execute intentions to take various steps that result in good health. We can be legitimately blamed for feeling mean-spirited, being disposed to lie frequently, reacting impatiently, forgetting a parent's birthday, failing to notice a friend is sad, and so on.<sup>23</sup> At best we have only indirect control over these things, because we accomplish them over a period of time by carrying out intentions to do other things. But this does not make us any less responsible for them.

## 4 The reason-responsiveness view

Given that direct intention-based control fails to satisfy the control condition for legitimate doxastic blame, we might be tempted to abandon intentions altogether and appeal to something else. In this spirit, advocates of what I will call the *Reason-Responsiveness View* argue that we have *reason-based* control over our beliefs.<sup>24</sup> An agent has *reason-based control* over her beliefs just in case she can form and maintain her beliefs by responding to reasons in the relevant way.

This sort of view has been motivated in several different ways. For example, [Ryan \(2003\)](#) and [Steup \(2008\)](#) each argue that figuring out what to believe is relevantly similar to figuring out what to do. In both types of deliberation, an agent considers and responds to various reasons—in one case, reasons to believe, and in the other, reasons to act. Ryan and Steup independently argue that if an agent is held responsible for her *actions* partly in virtue of her responsiveness to reasons, then agents can also be held responsible for their *beliefs* partly in virtue of their responsiveness to reasons.

Additionally, [Smith \(2005\)](#) argues that since we can be legitimately blamed for things over which we lack intention-based control, the best explanation of our control appeals to the fact that we have the capacity to respond to reasons. She holds that because our beliefs are responsive to reasons, they reflect our rational judgments, and hence we can be held responsible for them, even if we did not directly intend them.

### 4.1 Do we have reason-responsiveness?

I submit that whether or not an agent has reason-responsiveness depends on what it means to respond to reasons. So, first consider three different ways in which an individual might respond to information.<sup>25</sup> At the *first* and most basic level of information processing, an individual exhibits mere *sensitivity to stimuli*, like a plant that bends in response to sunlight or a litmus paper that turns red in response to acid. At the *second*

<sup>23</sup> See [Smith \(2005\)](#) for additional interesting examples of things for which we hold people responsible, but over which we lack direct intention-based control (which Smith calls “direct volitional control”).

<sup>24</sup> Supporters of various versions of the Reason-Responsiveness View include [Ryan \(2003\)](#) and [Steup \(2008, 2011, 2012\)](#), both of whom label their views “doxastic compatibilism.” See also [Hieronymi \(2006\)](#) and (2008, pp. 362–363), as well as [Smith \(2000, pp. 40–46\)](#) and (2005, pp. 236–271), who take their inspiration from [Scanlon \(1998\)](#) and [Korsgaard \(1996\)](#). Additional supporters include [Shah \(2002, p. 443\)](#) and [Weatherston \(2008, p. 546\)](#).

<sup>25</sup> I borrow here from Kornblith's discussion in (2012, pp. 50–53).

level, an individual additionally exhibits *sensitivity to reasons* and can engage in first-order reasoning. Creatures with this capacity—small children and various non-human animals—can modify their behavior and attitudes in response to reasons. They can integrate new information with stored information, thereby learning from new experiences. At the *third* level, an individual additionally exhibits *appreciation of reasons* and can engage in *second*-order reasoning. Creatures with the capacity to appreciate reasons—e.g. adult humans—can recognize their reasons *as* reasons and modify their behavior in response to this more complex recognition. Following Tyler Burge, such capacity is necessary for critical reflection.<sup>26</sup>

## 4.2 Is reason-responsiveness sufficient for doxastic control?

Obviously, advocates of the Reason-Responsiveness View appeal to *reason-responsiveness*, rather than non-rational stimuli, in order to satisfy the control condition for legitimate doxastic blame. This means that they must appeal either to sensitivity to reasons or to appreciation of reasons. However, I will argue in the next two sections that neither is sufficient for satisfying the control condition for legitimate doxastic blame.<sup>27</sup>

### 4.2.1 *Sensitivity to reasons*

First, consider agents who are sensitive to their reasons but can't recognize their reasons as such. If this capacity were sufficient to satisfy the control condition for legitimate doxastic blame, then we would expect that individuals who are sensitive to reasons, but who cannot appreciate them, could legitimately be blamed for their beliefs. But these are precisely the sort of individuals that we do *not* consider legitimate candidates for doxastic blame. As mentioned earlier, animals and young children, as well as babies, are the sorts of beings that are sensitive to evidence in the formation and maintenance of their beliefs. But we don't consider it legitimate to blame any of these individuals for their beliefs. And this seems right.

Indeed there are a variety of things that small children and animals believe on the basis of very poor grounds, for which we do not blame them. Children sometimes believe that the live actor they're hugging at Disneyland really *is* Mickey Mouse, that their stuffed animals come to life behind their backs, that the people they see on TV live inside the television, that monsters live in their closets, and so on.<sup>28</sup> Despite the

<sup>26</sup> Burge (1996, pp. 98–101).

<sup>27</sup> My thinking about reason-responsiveness with respect to believing has been greatly influenced by Peels (2013), in which he argues against six different versions of the view that agents can legitimately be subject to doxastic obligations even though they lack voluntary control over their beliefs. The conclusion that Peels draws is that we should endorse the view that we have some form of indirect intention-based control over our beliefs. I agree, but, as it will become clear, I go further in arguing that this control must take the specific form of reflection.

<sup>28</sup> Note that some children genuinely believe these things and aren't merely engaging in a form of pretense, even when no one tells them that they're true, and even when they don't watch movies or read books that would plant them in their heads.

fact that these beliefs are unjustified, we don't consider it legitimate to blame children for them. We certainly correct them, and try to show them the reasons for thinking otherwise. But if they fail to update their beliefs, we do not resent them or become indignant. Likewise for my dog when he continues to believe that my friend, who visits my house frequently without incident, is a threat. While I might well be annoyed, it's not legitimate for me to blame him. All of these cases indicate that the mere capacity for sensitivity to evidence, which children and dogs exhibit, is not sufficient for the sort of doxastic control that is required for legitimate doxastic blame.

The case is even stronger when we recall the nature of blame. Blame represents an agent's wrong belief in such a way that the agent can acknowledge her failure to respond to good epistemic reasons and can change her belief *in light of* this acknowledgement. In order for doxastic blame to be legitimate, it must be able, in principle, to play this causal role. But individuals that can only process information at the level of sensitivity to reasons cannot acknowledge that they believe on poor grounds. So they cannot modify their beliefs in response to this acknowledgement. Thus resenting a child for believing incorrectly or refusing to pet a dog as a form of protest for its belief does not have any point. Since reason-responsiveness in the form of mere sensitivity to reasons does not enable one to comply with the normative demand of doxastic blame, mere sensitivity to reasons is not sufficient to satisfy the control condition for legitimate doxastic blame.

#### 4.2.2 Appreciation of reasons

Now consider the capacity to appreciate reasons and not merely be sensitive to them. Consideration of our responses to the beliefs of children and animals indicates that in order for doxastic blame to be legitimate, the recipient of blame must be capable of responding to blame. To do this the agent must be able to influence her capacity for sensitivity to reasons in the right sort of way.<sup>29</sup> It's reasonable to think that this influence comes in the form of the ability to critically reflect, since this is what distinguishes children and animals from rational adults.

Is the capacity for critical reflection sufficient for satisfying the control condition for legitimate doxastic blame? I submit that, in order to answer this question, we must specify whether or not critical reflection is something an individual can engage in as the result of an intention. If something is done as a result of an intention, then, for simplicity, I'll call it *active*, and if something is not done as a result of an intention, I'll call it *passive*.

Consider what it would be for an agent to engage in passive reflection. Anyone who's ever tried to figure out a solution to a complicated problem has likely experienced finding themselves thinking about what they're going to eat for dinner that night, or suddenly remembering the phone call they're supposed to make, or wondering if they're going to hear back from that job interview two weeks ago. Thoughts and memories can pop into our head uninvited, or we can intend to think about them, intend to direct our attention toward them, intend to try to remember them. An agent engaged

<sup>29</sup> See Peels (2013, pp. 15–16) for an additional argument that such influence is required for agents to be held responsible for their beliefs.

in reflection passively would be in a similar situation with respect to the various mental events that compose reflection. She would think about various things without having actively directed her thoughts to them. She would find herself wondering whether a certain reason supports believing a proposition, though she didn't intend to direct her attention there. She would suddenly remember a piece of evidence that she'd previously forgotten, and so forth.

I contend that the capacity for passive reflection is not sufficient to satisfy the control condition for legitimate doxastic blame. One might argue that it's never the case that we engage in critical reflection passively, perhaps because reflection, like long division, necessarily involves our intentions. I suspect this is right, but the point is that *if* one could engage in critical reflection only passively, this ability would not suffice for the doxastic control required for legitimate doxastic blame.

The capacity for active reflection, however, *is* sufficient to satisfy the control condition for legitimate doxastic blame. If an agent can actively engage in reflection, then she has control over doing various mental actions that influence her beliefs. Being able to do this enables an individual to respond to any blame incurred by her beliefs. If she is legitimately blamed for believing that  $p$ , then in response to the blame, she can carry out various intentions to reflect on her reasons for believing  $p$ . Doing so helps her recognize or understand that believing that  $p$  amounts to epistemic failure.

So, I pose a dilemma: either the capacity for reflection can be exercised passively or it cannot. If it *can* be exercised passively, then we can imagine someone who can at best only passively engage in reflection, and when we do, we find that intuitively this person cannot legitimately be blamed for her resulting beliefs. On the other hand, if reflection *cannot* be exercised passively—if it can only be exercised actively—then I submit that part of the reason why the capacity for active reflection satisfies the control condition for legitimate doxastic blame is that it involves our intentions. Without the ability to reflect as a result of an intention, we cannot be legitimately blamed for our beliefs.

The upshot is that a crucial part of the characterization of our doxastic control involves our intentions: we must have intention-based control over whether we engage in the activity of critical reflection, or more specifically, the activities that compose reflection. Of course, this need not be *direct* intention-based control as I pointed out in Sect. 3.2. And to be clear, it's not that an agent must engage in active reflection about a proposition  $p$  in order to be legitimately blamed for believing that  $p$ . Rather, it's that an agent must be *able* to engage in active reflection about  $p$  in order to be blameworthy for believing that  $p$ .

Now let's consider whether active reflection is a form of reason-responsiveness. In other words, if we appeal to the capacity for active reflection to satisfy the control condition for legitimate doxastic blame, then is it fair to say that we have doxastic control in virtue of our ability to respond to reasons?

I submit that it is not. To be able to actively reflect, one must be able to respond to reasons—in the form of both sensitivity to reasons and sensitivity to appreciation of reasons. However, such responsiveness is insufficient on its own to constitute the control required for legitimate doxastic blame. In order to satisfy the control condition, one must additionally be able to execute intentions to carry out the mental actions that constitute reflection. What helps explain the legitimacy of doxastic blame is not our

ability to respond to reasons per se, but rather our ability to perform mental actions that influence the process of responding to reasons. These actions make a difference to whether we believe a proposition. Therefore, it would be misleading to characterize our doxastic control as responsiveness to reasons. We ought instead to characterize it in terms of our capacity for active reflection.

## 5 The reflective control view

We are now in a position to see the motivation for the *Reflective Control View* of doxastic control. In Sect. 3 I argued that direct intention-based control over our beliefs is *not necessary* for satisfying the control condition for legitimate doxastic blame. In Sect. 4 I argued that reason-based control is *not sufficient* for satisfying the control condition. The Reflective Control View corrects the shortcomings of both of these views. According to this view, we have what I will call *indirect reflective control* over our beliefs:

An individual has *indirect reflective control* over whether she believes that  $p$  if and only if she can actively engage in critical reflection that causally influences whether or not she holds the belief that  $p$ .

The sort of reflection involved in indirect reflective control is second-order in the sense that it involves conceptualizing relations of evidential support between various propositions or between one's attitudes towards those propositions. The mental actions that compose reflection include asking ourselves various questions and directing our attention in various ways. One might ask what one's reasons for believing  $p$  are, whether those reasons are good, whether they're sufficient for belief, whether one has any new reasons for or against  $p$ , whether one has interpreted one's evidence correctly, whether one is biased in any significant way, and so forth. One might direct one's attention toward a certain class of reasons, such as relevant memories, or away from certain reasons. One might try to compare or contrast the case at hand with other similar cases. Trying to remember various reasons, synthesizing them, reconstructing an argument for a proposition—all these are examples of the various mental actions that are part of critical reflection. And all of these things causally influence whether a person believes a proposition.

Reflection on propositions about which we haven't yet formed any attitude causally influences our belief *formation*; and reflection on propositions we already believe causally influences our belief *maintenance*. I don't propose a specific account of causal influence because the proposal is intended to be neutral between competing theories of causal influence.

According to the Reflective Control View, indirect reflective control is both necessary *and* sufficient for satisfying the control condition for legitimate doxastic blame. Thus, an agent can be legitimately blamed for her belief that  $p$  only if she has indirect reflective control over whether she believes that  $p$ . Note that one has indirect reflective control just in case one has the *capacity* for the relevant reflection, regardless of whether one in fact reflects. To have control one need not exercise the capacity; one simply must have it. By way of analogy, I have indirect control over how much

I weigh because I have the capacity to carry out an intention to eat fewer calories, whether or not I actually do so. The Reflective Control View is the natural outcome of reflecting on the deficiencies of both the direct intention-based control view and the Reason-Responsiveness View, though it bears some similarity to each.

Before considering those similarities, first note a difference between the Reflective Control View and each of the other two views: the Reflective Control View advocates a form of *indirect* control. Our control over what we believe is mediated by our ability to execute various intentions to do the activities that compose reflection. On the other two views, our doxastic control is *direct*: one's sensitivity to reasons or one's appreciation of reasons causally affects one's beliefs directly, as does one's intention to believe a particular proposition. Recall that appealing to indirect control is not a problem for explaining doxastic blame since there are many things for which we can be legitimately blamed, even though we only have indirect control over them.

Though *direct* intention-based control is not necessary for blame, one might wonder whether the control I'm advocating—indirect reflective control—is a form of *indirect* intention-based control. After all, the Reflective Control View is similar to intention-based views of control in that both views appeal to intentions to explain doxastic control. However, intention-based doxastic control views advocate control over our beliefs by carrying out intentions to believe specific propositions, whether directly or indirectly; we control whether we believe that  $p$  via whether we intend to believe that  $p$ . But on the Reflective Control View, we have control over whether we believe that  $p$  via whether we intend to engage in various reflective activities which causally influence whether we believe that  $p$ . So the controlling intention is not an intention to believe a particular proposition—indeed, one need not have any such intention. Instead, the controlling intention is an intention to engage in reflection. More generally, this is a contrast between intention-based control, whereby we control whether we  $\phi$  by whether we intend to  $\phi$ , and what I will call *influence-based control*, whereby we control whether we  $\phi$  by whether we intend to  $\psi$  which then causally influences whether we  $\phi$ . The Reflective Control View appeals to a form of influence-based control rather than intention-based control. Moreover, the specific kind of influence is reflective influence.

The Reflective Control View is similar to the Reason-Responsiveness View insofar as both views prioritize the ability to respond to reasons. The difference is that on the Reflective Control View, reason-responsiveness is necessary *but not* sufficient for satisfying the control condition for legitimate doxastic blame. The Reflective Control View adds that we must have the ability to actively engage in critical reflection about our beliefs. Thus, the kind of control that satisfies the control condition for legitimate doxastic blame comes in the form of control over certain actions—those actions involved in reflection.

## 6 Defending the reflective control view

At this point, I want to address two objections to the Reflective Control View. According to the first objection, indirect reflective control is not sufficient to satisfy the control condition for legitimate doxastic blame, and according to the second, it's not necessary.

## 6.1 The objection from sufficiency

First, we consider the objection that indirect reflective control is not sufficient to satisfy the control condition. There are two forms of this objection. According to one, indirect reflective control is not sufficient because reflection, whether active or not, often lacks causal influence—a large number of our beliefs are such that reflection fails to causally influence them. According to the other form of the objection, when reflection does influence our beliefs, it usually has a negative influence—e.g. it causes us to believe *out* of accord with our evidence. The thought is that we often reflect poorly: we maintain unjustified beliefs because we irrationally conclude from reflection that the belief is justified, and we abandon justified beliefs because we irrationally conclude from reflection that the belief is unjustified. But the sort of control that satisfies the control condition both causally influences our beliefs and helps us believe in accordance with epistemic norms so that we can avoid incurring doxastic blame. Therefore, active reflection fails to satisfy the control condition for legitimate doxastic blame.

My first response to these objections is to concede that in order for the capacity for active reflection to satisfy the control condition for legitimate doxastic blame, reflection must have a positive epistemic influence. If it were the case that engaging in reflection always negatively influenced our beliefs, then we would be epistemically better off refraining from reflection, given that it would hinder us from believing in accordance with epistemic norms. If we cannot engage in reflection *well*—perhaps because we're just not wired in the right way—then we cannot be said to have the kind of doxastic control that makes us legitimate candidates for doxastic blame. So we should add the following modification to the notion of indirect reflective control: an agent has indirect reflective control if and only if she has the capacity to actively engage in reflection that has a *positive* epistemic influence on her beliefs. Indirect reflective control, understood in this way, *is* sufficient, in principle, for satisfying the control condition for legitimate doxastic blame.

But in light of the ways in which reflection can go wrong, one might wonder how often we have indirect reflective control, or whether we ever have it. Note that it's an empirical question whether, and if so how often, we have the capacity to reflect in a way that positively influences our beliefs. We might learn from empirical research that we have indirect reflective control over some of our beliefs but not all of them. Indeed, I think this is what we should expect to find: we don't always have the capacity to positively influence our beliefs via reflection. But this also fits with our doxastic practices, given that we don't think that agents can be legitimately blamed for all of their beliefs but only some of them.

Suppose, however, that empirical research reveals the more sweeping conclusion that reflection *never* influences our beliefs positively, either because it fails to exert any causal influence at all or because it always exerts a negative influence. In other words, suppose that we never have indirect reflective control. While this supposition would not falsify the claim that indirect reflective control is sufficient to satisfy the control condition for legitimate blame, I've also claimed that indirect reflective control is necessary for satisfying it. So if empirical findings reveal that we never have indirect reflective control, then on my view we'd have to conclude that it's never legitimate to blame people for their beliefs.

While this might seem a surprising result, ultimately, it's an empirical question whether we can ever be blamed legitimately for our beliefs. While I consider our practice of doxastic blame legitimate until proven otherwise, I'm open to the possibility of learning from the empirical research that doxastic blame is legitimate in far fewer cases than our current practice assumes, perhaps even that it's *never* legitimate. However, I don't think that our empirical research to date shows us that doxastic blame is never legitimate, because I don't think that it conclusively shows that we never have indirect reflective control. When we examine the findings in detail, I think we find that cases in which agents lack indirect reflective control (i.e. in which the agent cannot engage in reflection with a positive causal influence) are also cases in which it's reasonable to think that agents cannot legitimately be blamed for their beliefs. To see this let's look at some of the research.

Kornblith (2012) supports the claim that reflection often lacks causal influence by citing various empirical studies that indicate that our beliefs are often influenced in powerful ways by non-evidential factors of which we are not even conscious, such as various smells or visual input. For example, people tend to believe that objects further to the right are superior to those on the left, though product placement has nothing to do with the quality of a product; people's beliefs about others are influenced by exposure to words, like "rude" and "polite"; and beliefs about politicians are influenced by the random colors on their posters (2012, pp. 21–22).<sup>30</sup> The thought is that these factors are omnipresent, thus affecting a wide variety of our beliefs. Since these factors are below the level of consciousness, we can't correct for them via reflection; thus, engaging in reflection on such matters doesn't causally influence what we believe. As Kornblith puts it, "Asking subjects to introspect more carefully, or think longer and harder about the sources of their beliefs, is entirely useless in many of these cases" (2012, p. 23).

If we grant that these studies show that agents cannot reflect to positive causal result in the cases described, then on my view we must conclude that agents cannot be legitimately blamed for their beliefs in such cases. But I think such a conclusion matches our intuitions; the conclusion seems reasonable when we examine particular cases in detail. For example, suppose Kate incorrectly believes that Colgate is a more effective toothpaste than Crest, but unbeknownst to her she's influenced to believe this by the fact that Colgate toothpastes are located to the right of Crest ones in the shopping aisle. Assume that if she were to reflect, she would not be able to identify this influence. In this case, it seems unintuitive to think that Kate is legitimately blameworthy for her belief. In other words, though we grant that Kate lacks indirect reflective control over her belief about the more effective toothpaste, it's also the case that she cannot be legitimately blamed for the belief. This verdict seems reasonable; it matches our practices.

Now consider Kornblith's support of the claim that reflection often results in negative causal influence. He argues that in a large class of cases, engaging in reflection is "a route to little more than self-congratulation": it allows us to irrationally strengthen our confidence in the unjustified conclusions about ourselves and others that we formed apart from critical reflection in the first place (2012, pp. 3, 25). Referring to the large

---

<sup>30</sup> See Kornblith (1989) for more examples.



body of psychological literature on overconfidence, Kornblith also points out that we often overestimate the accuracy of our second-order judgments, such as judgments about the import of our evidence (p. 25). And when reflecting on evidence for beliefs that we already hold, we are strongly disposed to scrutinize evidence against them but typically take evidence in favor of them at face value, which also leads to inaccurate beliefs (p. 25). Kornblith concludes that reflection “by and large” is not an improvement over our first-order processes; it does not guard us from cognitive errors; and it does not aid overall in the project of cognitive self-improvement (p. 26).

Again, if we grant that the studies show that agents cannot reflect to positive causal result in the cases described, then on my view we must conclude that agents cannot be legitimately blamed for their beliefs in such cases. But again, I think this is the correct result. Do we really think it’s legitimate to blame someone for an incorrect psychological belief about their abilities if no amount of reflection could possibly dislodge that belief? I think not. In such a case, it’s unfortunate that the person has an incorrect belief, but since she lacks the capacity for causally influential reflection, she cannot acknowledge that the belief is incorrect and thus that she should give it up. Given that she would not be able to respond to the demand inherent in the blame, it would not be legitimate to blame her for the belief.

Lastly, I want to end by highlighting that though I’ve explored the consequences of assuming that reflection lacks positive causal influence, it’s worth pointing out that the empirical research is by no means conclusive with respect to the causal impact of reflection. While Kornblith cites various studies according to which engaging in reflection either fails to influence our beliefs or influences them negatively, there are also various empirical studies that show that engaging in reflection can influence our beliefs positively. In other words, there’s evidence that reflection does influence our beliefs to *better* accord with epistemic norms.

Psychologists Roy F. Baumeister, E. J. Masicampo, and Kathleen D. Vohs summarize a number of these findings in their (2011) review. Many of the studies they review indicate that conscious deliberation improves an agent’s ability to reason well and thus, to perform well. Subjects engaging in problem solving, who were told that they would later be asked to explain their reasoning, were more successful at solving the problems (p. 341).<sup>31</sup> Various subjects asked to verbalize their reasoning while performing certain tasks performed those tasks better (p. 341).<sup>32</sup>

There are also various studies that show that many mental processes that are assumed to be automatic can be altered by conscious control (p. 350). For example, according to the victim effect, people tend to donate more when they encounter specific identifiable victims in need than when they encounter abstract statistics describing the need. But in one study, this effect was counteracted by priming subjects to deliberate about their decision to donate (p. 350).<sup>33</sup> Another study indicates that the causal effect of

---

<sup>31</sup> See the studies reported by DeWall et al. (2008).

<sup>32</sup> For one such study, see Gagne and Smith (1962).

<sup>33</sup> See Small et al. (2007).

pro-smoking and anti-smoking cues is mitigated when a person is able to consciously reflect on them (p. 351).<sup>34</sup>

Given these studies and others, the empirical case against the positive causal influence of reflection is not conclusive. Indeed, we have empirical reason to think that engaging in reflection does influence a person to form and maintain justified beliefs. Thus, I maintain that we do have indirect reflective control over many of our beliefs, and this control helps legitimate our practice of doxastic blame.

## 6.2 The objection from necessity

We now consider the objection that indirect reflective control is not necessary for satisfying the control condition for legitimate doxastic blame. Imagine an agent who cannot engage in reflection at all, but further imagine that she has access to a special doxastic hat that will ensure that she believes in accordance with her evidence whenever she's wearing it. Assume that she knows what happens when she wears the hat, and that she can choose to wear it at any time. She thus has indirect control over her beliefs in the form of the ability to actively put on the hat. And moreover, it seems that if she were obligated to believe in accordance with her evidence, then given her access to, and knowledge of, the hat, if she refused to wear it, we could blame her for any of her beliefs that don't accord with her evidence.

According to this objection, reflection, whether active or not, is not necessary for explaining the legitimacy of doxastic blame because various other non-reflective activities, such as wearing a doxastic hat, will do the job. Indeed there are many non-reflective activities that indirectly influence our beliefs.<sup>35</sup> At a very basic level, I can influence whether I believe certain things by directing my eyes various places, thus generating visual evidence on which I base my beliefs. Additionally, I can influence my beliefs by choosing whether to stay up late watching Netflix instead of studying my *Book of Common Errors in Reasoning*. Even activities like gathering more evidence, seeking out a second opinion, and reading relevant books about the topic in question fall into the category of non-reflective activities. Alston cites these as examples of ways we can indirectly influence our beliefs (1988, pp. 278–279). Insofar as Alston's notion of voluntary influence includes these activities, my preferred notion of indirect reflective control is a narrower form of influence-based control.

Olson (2015) also appeals to these sorts of activities in his defense of epistemic agency. Though it's not Olson's goal to defend the legitimacy of doxastic blame, he argues that his notion of epistemic agency is needed to account for the legitimacy of various normative epistemic notions. According to Olson, epistemic agency is the agency we have over our belief-forming practices which, when exercised, directly affects the way in which we form beliefs and indirectly affects the beliefs we form (2015, p. 449). While Olson agrees that reflection is an important part of epistemic agency—"Reflection and deliberation, e.g., are types of practices that can influence belief formation and are things we can control" (2015, p. 452)—he does not claim

<sup>34</sup> See Westling et al. (2006).

<sup>35</sup> Olson (2015) nicely outlines many of them.

that the capacity for reflection is necessary for epistemic agency. Epistemic agency on Olson's view is practical agency—the agency employed in executing plans over time—directed toward epistemic matters.

However, while these non-reflective activities are things we *can* voluntarily do, and while they *do* have a causal impact on our beliefs, they are not the right type of thing to appeal to in order to satisfy the control condition for legitimate doxastic blame. By themselves they are not sufficient for the control required for legitimate blame. We can see this by noting that even small children and animals can voluntarily engage in such non-reflective activities, yet they are not candidates for legitimate blame.

More importantly, there is principled reason why legitimate doxastic blame requires indirect *reflective* control as opposed to indirect non-reflective control. Whatever satisfies the control condition for legitimate blame must be the sort of thing that enables us to respond to blame in the right way. Responding to blame in the right way requires responding to the demand that others make on us when they blame us. The demand is to acknowledge one's failure to believe correctly given one's reasons. Doing this requires the capacity for reflection. If I lack indirect *reflective* control—i.e. the capacity to engage in active reflection that makes a positive epistemic difference to my beliefs—then I can't respond to the demand that you place on me when you blame me. I can't acknowledge that my belief is incorrect given my reasons, because I can't recognize my reasons as good or bad. Similarly, if there is some action I could perform that would ensure that I believe correctly, I cannot be expected to perform it if I lack the capacity for active reflection. Without the capacity for active reflection, I cannot actively recognize as a reason to perform the action that it ensures correct beliefs.<sup>36</sup>

So, without the capacity for reflection, I can't be blamed for having an unjustified belief as result of refusing to study my logic book. In order to be legitimately blamed, I need to be able to recognize the connection between not studying the book and not believing correctly, and I need to be able to actively do what enables me to recognize that connection, i.e. reflect. Likewise, I cannot be blamed for having an unjustified belief that *p* even if I am capable of gathering more evidence unless I can also recognize my need to gather more evidence as a reason to refrain from believing that *p*. But I can't do this if I lack the capacity for active reflection. Even for an individual who has access to the doxastic hat, in order to be legitimately blamed for her belief given her failure to wear the hat, the individual would need the capacity for reflection in order to understand why she would be expected to wear the hat in the first place. Without it, she would not be able to grasp that since it's good to have justified beliefs, she therefore has a reason to wear the hat.

<sup>36</sup> In keeping with Burge (1996), I hold that reflection is necessary for agents to be able to recognize their reasons as such. While much can be said about the relationship between reflection and recognition of reasons, in brief, here's how I'm thinking of it: The reasons that an agent has depend on that agent's mental states. So in order to recognize one's reasons, one must be able to recognize one's mental states, which is something an agent can only do via reflection. An agent might learn of her own mental states in some other way—by testimony, for example. But learning via testimony, at best, produces a belief that falls short of recognition. Recognizing one's reason as a reason does not amount to merely believing that one's reason is a reason. It involves a kind of understanding of one's own mental states which is something one cannot do apart from reflection. Thanks to Declan Smithies for helpful discussion and to an anonymous reviewer for pressing me to clarify this.

In order to demonstrate that the capacity to engage in non-reflective activities is sufficient in the absence of the capacity for reflective ones, we would need to imagine an agent who lacks the capacity for active reflection and yet can still be legitimately blamed for her belief. But when we consider our practices of doxastic blame, we find that individuals who lack the capacity for active reflection are precisely the sort of individuals that we do not blame. It's not legitimate to blame children and animals for their beliefs; it's not even legitimate to blame an adult with incorrect beliefs if that individual lacks the ability to influence her beliefs via reflection. These agents cannot be legitimately blamed because they are incapable of responding to the demand inherent in blame. The capacity for active reflection is required to respond to this demand. So I maintain that the capacity for active reflection *is* necessary for satisfying the control condition for legitimate doxastic blame.

If the capacity to engage in active reflection that makes a positive difference to one's beliefs—i.e. indirect reflective control—is required in this way for legitimate doxastic blame, one might wonder whether such reflective control is also required for legitimately blaming individuals for their health and character traits. Recall from Sect. 3.2 that I pointed out that our control over such things is indirect. Here I think we find that indirect reflective control *is* required, though the influence of active reflection is even more indirect than its influence with respect to beliefs and other doxastic attitudes. Indeed, it's plausible that the same is true of various attitudes besides belief, such as implicit biases and prejudices, though a full defense of this must be reserved for another paper. For now, I simply note the attractiveness of appealing to indirect reflective control to satisfy the control condition for legitimate blame for a variety of things over which we lack direct intention-based control.

## 7 Conclusion

I have argued that we do have the sort of control over our beliefs required for legitimate doxastic blame. I distinguished between three types of doxastic control: intention-based, reason-based, and influence-based control. I argued first that though we lack direct intention-based control over our beliefs, such control is not necessary for legitimate doxastic blame. Second, reason-based control is insufficient to satisfy the control condition for legitimate doxastic blame if we interpret reason-responsiveness as mere sensitivity to reasons. We might consider reason-responsiveness to amount to critical reflection. But in that case we must distinguish passive from active reflection. Only active reflection is capable of satisfying the control condition. Furthermore, I argued that thinking of active reflection as a form of reason-responsiveness masks the fact that our doxastic control is in fact a form of influence-based reflective control—it is the capacity to execute intentions to engage in reflection that causally influences our beliefs in positive epistemic ways. This capacity is both necessary and sufficient for satisfying the control condition for legitimate doxastic blame.

**Acknowledgements** I presented earlier versions of this paper at the 2014 Pacific APA and the 2015 Alabama Philosophical Society; thanks to those in attendance for their questions and objections. Thanks to Ryan Bond, Blake McAllister, Caroline Paddock, Hope Sample, and Keren Wilson for productive exchanges about the ideas here; and thanks to Matt Frise, Allan Hazlett, Dan Howard-Snyder, Daniel Immerman, Abe

Roth, Meg Schmitt, Tim Schroeder, and two anonymous referees for detailed comments. Special thanks to Bradley Rettler for several helpful discussions and to Declan Smithies for many invaluable conversations and extensive comments on multiple drafts.

## References

- Alston, W. (1985). Concepts of Epistemic Justification. In W. Alston (Ed.), *Epistemic justification: Essays in the theory of knowledge* (pp. 81–114). Ithaca: Cornell University Press.
- Alston, W. (1988). The deontological conception of epistemic justification. *Philosophical Perspectives*, 2, 257–299.
- Baumeister, R., Masicampo, E. J., & Vohs, K. D. (2011). Do conscious thoughts cause behavior? *Annual Review of Psychology*, 62, 331–361.
- Buckareff, A. (2014). Deciding to Believe Redux. In J. Matheson & R. Vitz (Eds.), *The ethics of belief: Individual and social* (pp. 33–50). Oxford: Oxford University Press.
- Burge, T. (1996). Our entitlement to self-knowledge. *Proceedings of the Aristotelian Society*, 96, 91–116.
- Clifford, W. K. (1999). The Ethics of Belief. In T. Madigan (Ed.), *The ethics of belief and other essays* (pp. 70–96). Amherst: Prometheus Books.
- DeWall, N. C., Baumeister, R., & Masicampo, E. J. (2008). Evidence that logical reasoning depends on conscious processing. *Consciousness and Cognition*, 17, 628–645.
- Feldman, R. (2000). The ethics of belief. *Philosophy and Phenomenological Research*, 60, 667–695.
- Feldman, R. (2001). Voluntary Belief and Epistemic Evaluation. In M. Steup (Ed.), *Knowledge, truth, and duty* (pp. 77–92). Oxford: Oxford University Press. chapter 5.
- Gagne, R. M., & Smith, E. C. (1962). A study of the effects of verbalization on problem solving. *Journal of Experimental Psychology*, 63, 12–18.
- Hieronymi, P. (2004). The force and fairness of blame. *Philosophical Perspectives*, 18, 115–148.
- Hieronymi, P. (2005). The wrong king of reason. *Journal of Philosophy*, 102, 437–457.
- Hieronymi, P. (2006). Controlling attitudes. *Pacific Philosophical Quarterly*, 87, 45–76.
- Hieronymi, P. (2008). Responsibility for believing. *Synthese*, 161, 357–373.
- Hieronymi, P. (2009). Believing at will. *Canadian Journal of Philosophy*, 35, 149–187.
- Kelly, T. (2002). The rationality of belief and some other propositional attitudes. *Philosophical Studies*, 110, 163–196.
- Kornblith, H. (1989). Introspection and misdirection. *Australasian Journal of Philosophy*, 67, 410–422.
- Kornblith, H. (2012). *On reflection*. Oxford: Oxford University Press.
- Korsgaard, C. (1996). Morality as Freedom. In C. Korsgaard (Ed.), *Creating the kingdom of ends*. Cambridge: Cambridge University Press.
- McKenna, M. (1998). The Limits of evil and the role of moral address: A defense of Strawsonian compatibilism. *Journal of Ethics*, 2(2), 123–142.
- Olson, D. (2015). A case for epistemic agency. *Logos and Episteme*, 6, 449–474.
- Peels, R. (2013). Doxastic compatibilism. *Philosophy and Phenomenological Research*, 89, 679–702.
- Ryan, S. (2003). Doxastic compatibilism and the ethics of belief. *Philosophical Studies*, 114, 47–49.
- Scanlon, T. M. (1988). The significance of choice. *Tanner Lectures on Human Values*, 8, 149–216.
- Scanlon, T. M. (1998). *What we owe to each other*. Cambridge: Harvard University Press.
- Scanlon, T. M. (2008). *Moral dimensions: Permissibility, meaning, and blame*. Harvard: Harvard University Press.
- Setiya, K. (2008). Believing at will. *Midwest Studies in Philosophy*, 32, 36–52.
- Shah, N. (2002). Clearing space for doxastic voluntarism. *The Monist*, 85, 436–445.
- Shah, N. (2006). A new argument for evidentialism. *The Philosophical Quarterly*, 56, 481–498.
- Sher, G. (2006). *In praise of blame*. Oxford: Oxford University Press.
- Small, D. A., Loewenstein, G., & Slovic, P. (2007). Sympathy and callousness: The impact of deliberative thought on donations to identifiable and statistical victims. *Organizational Behavior and Human Decision Processes*, 102, 143–153.
- Smith, A. (2000). Identification and responsibility. In T. van den Beld (Ed.), *Moral responsibility and ontology* (pp. 233–246). Dordrecht: Kluwer Academic Publishers.
- Smith, A. (2005). Responsibility for attitudes: Activity and passivity in mental life. *Ethics*, 115, 236–271.
- Smith, A. (2008). Control, responsibility, and moral assessment. *Philosophical Studies*, 138, 367–392.

- Smith, A. (2013). Moral blame and moral protest. In D. J. Coates, & N. A. Tognazzini (Eds.), *Blame: Its nature and norms* (pp. 27–48). Oxford: Oxford University Press.
- Steup, M. (2008). Doxastic freedom. *Synthese*, *161*, 375–392.
- Steup, M. (2011). Belief, voluntariness and intentionality. *Dialectica*, *65*, 537–559.
- Steup, M. (2012). Belief, control and intentionality. *Synthese*, *188*, 145–163.
- Strawson, P.F. (Ed.). (1974). Freedom and resentment. In *Freedom and resentment and other essays*, (pp. 1–25). London: Methuen.
- van Woudenberg, R. (2009). Responsible belief and our social institutions. *Philosophy*, *84*, 47–73.
- Wallace, R. J. (1994). *Responsibility and the moral sentiments*. Cambridge: Harvard University Press.
- Wallace, R. J. (2007). Reasons, relations, and commands: Reflections on Darwall. *Ethics*, *118*, 24–36.
- Wallace, R. J. (2010). Hypocrisy, moral address, and the equal standing of persons. *Philosophy and Public Affairs*, *38*, 307–341.
- Wallace, R. J. (2011). Dispassionate opprobrium: On blame and the reactive sentiments. In R. J. Wallace, R. Kamar, & S. Freeman (Eds.), *Reasons and recognition: Essays on the philosophy of T.M. Scanlon* (pp. 348–372). New York: Oxford University Press.
- Watson, G. (Ed.). (2003). *Free will* (2nd ed.). Oxford: Oxford University Press.
- Weatherston, B. (2008). Deontology and Descartes' demon. *Journal of Philosophy*, *105*, 540–569.
- Westling, E., Mann, T., & Ward, A. (2006). Self-control of smoking: When does narrow attention help? *Journal of Applied Social Psychology*, *36*, 2115–2133.
- Williams, B. (1973). *Deciding to believe* (pp. 136–151). Cambridge: Cambridge University Press.
- Wolf, S. (1980). Asymmetrical freedom. *The Journal of Philosophy*, *77*, 151–166.
- Wolf, S. (1990). *Freedom within reason*. New York: Oxford University Press.
- Wolf, S. (2011). Blame, Italian style. In R. J. Wallace, R. Kamar, & S. Freeman (Eds.), *Reasons and recognition: Essays on the philosophy of T.M. Scanlon* (pp. 131–213). New York: Oxford University Press.