

Against Doxastic Compatibilism

RIK PEELS

VU University Amsterdam

Abstract

William Alston has argued that the so-called deontological conception of epistemic justification, on which epistemic justification is to be spelled out in terms of blame, responsibility, and obligations, is untenable. The basic idea of the argument is that this conception is untenable because we lack voluntary control over our beliefs and, therefore, cannot have any obligations to hold certain beliefs. If this is convincing, however, the argument threatens the very idea of doxastic responsibility. For, how can we ever be responsible for our beliefs if we lack control over them? Several philosophers have argued that the idea that we bear responsibility for our beliefs can be saved, because absence of voluntary control over our beliefs is perfectly compatible with having obligations to hold particular beliefs. With others, I call this view ‘doxastic compatibilism’. It comes in two varieties. On the first variety, doxastic obligations do not require any kind of doxastic control whatsoever. I argue that this variety of doxastic compatibilism fails because it confuses doxastic *responsibility* with other closely related phenomena. On the second variety, doxastic obligations do not require voluntary doxastic control, but only compatibilist doxastic control (roughly, reason-responsiveness) and we do in fact have such control. I grant that we have such control, but also argue that having such control is insufficient for bearing doxastic responsibility. The plausibility of the examples put forward by doxastic compatibilists in support of the claim that it *is* sufficient for doxastic responsibility derives from the fact that in these examples, the subjects have control over factors that influence what they believe rather than control over those beliefs themselves.

1. Introduction

Several philosophers, most notably William Alston, have provided more or less the same argument against the so-called *deontological conception of epistemic justification*.¹ On this conception, epistemic justification should be understood in terms of obligations, permissions, requirements, responsibility, praise, and blame. The idea is roughly that some person S justifiedly believes that p iff in believing that p , S does not violate any obligations and is, therefore, blameless for believing that p . And, of course, *mutatis mutandis* the same is taken to be true for other doxastic attitudes, such as disbelief and withholding. Alston has argued that we lack voluntary control over our beliefs, but that such control is necessary for having doxastic obligations. In order to understand the thesis that we lack voluntary control over our beliefs, let me say a few words on control. By ‘voluntary control’ Alston means *intentional* and *libertarian* control. One has voluntary control over something if and only if one can choose to do it and one can choose not to do it or, in other words, if one can decide to do it and one can decide not to do it. Slightly more precisely, the idea seems to be that one has voluntary control over φ -ing if and only if one can φ as the result of an intention to φ and one can $\sim\varphi$ as the result of an intention to $\sim\varphi$.² Thus, Alston would say, in normal circumstances I have voluntary control over what I eat for breakfast, what I say to my colleagues, how much time I spend on my work, how I behave when I drive my car, and so on. For in all these cases there is something which I can decide to do and decide not to do. According to Alston, it is clear that we lack this kind of control over the vast majority of our beliefs: I cannot change my beliefs simply as the result of an intention to do so in the way I *can* choose to say something to a colleague or in the way I *can* choose to spend an hour on writing a letter. I believe that I had two slices of bread for breakfast, that I am now in my study, that it is a sunny day, and so on, but I have no idea how I could now abandon these beliefs. It seems that my beliefs are not under my voluntary control.

Furthermore, it seems that we can only have an obligation to believe something if we have such control (‘ought’ implies ‘can’). Hence, we do not have obligations to hold particular beliefs. But then being epistemically justified in believing a proposition cannot be a matter of meeting or, at least, not violating one’s doxastic obligations. Slightly more formally:

- (1) We cannot have obligations to believe particular propositions unless we have sufficient voluntary control over our beliefs.

¹ See, for instance, Alston 1989, 115-152; 2005, 58-60.

² Even more precisely, one can φ and one can $\sim\varphi$ as the result of an intention to φ or $\sim\varphi$ via a non-deviant causal chain.

(2) We do not have sufficient voluntary control over our beliefs.

(3) Hence, we do not have doxastic obligations.³

Now, if epistemic justification should be understood in terms of obligations to believe and we do not have such obligations, then the deontological conception of epistemic justification is in trouble. Just to be explicit, here is why. If being epistemically justified is supposed to be a matter of *meeting* one's doxastic obligations, then one would *never* be epistemically justified, since there would not be any such obligations. If being epistemically justified is supposed to be a matter of *not violating* one's doxastic obligations, there are two options. First, such a view could presume that there are doxastic obligations. It would then follow that we are *never* epistemically justified. Second, such a view could *not* presume that there are doxastic obligations. We would then *always* be epistemically justified. But any plausible theory of epistemic justification should allow that sometimes we believe justifiedly and sometimes we do not.

It is important for our purposes to note that this argument not only threatens the deontological conception of epistemic justification, but also the very idea that we can properly hold each other responsible for our beliefs. We believe that—unless special conditions hold—the racist's belief is blameworthy. And we hold each other responsible for our moral and political beliefs.⁴ The whole philosophical discussion about a viable ethics of belief—which we already find in the works of such philosophers as John Locke and William Clifford—presumes that we *are* indeed responsible for our beliefs. Sometimes we believe responsibly and sometimes we believe blameworthily or culpably. In this paper, I will not consider whether the deontological conception of epistemic justification is tenable. In fact, I will not be explicitly concerned with epistemic justification at all and focus entirely on responsibility for our beliefs.

Several strategies have been devised to save doxastic responsibility from the above argument. Some philosophers have claimed that, surprisingly, we *do* have some kind of voluntary control over our beliefs.⁵ Others have argued that we lack control over our beliefs, but that we nonetheless *influence* them via our control over all sorts of factors that make a difference to what we believe, factors such as evidence-gathering and working on our

³ For this argument, see Alston 1989, 115-136.

⁴ For linguistic evidence for this claim, see Van Woudenberg 2009.

⁵ That doxastic responsibility can at least partly be explained in terms of our indirect doxastic control has been claimed by Heil 1992, 51; Nottelmann 2007, 157-159; Price 1954, 16-21; Wolterstorff 2010, 62-85.

intellectual virtues and vices.⁶ However, it seems that the major strategy has been to grant that we lack voluntary control over our beliefs, but that we nonetheless have doxastic obligations. In other words, the idea is that having doxastic obligations and *not* having doxastic voluntary control are perfectly compatible with each other. I, therefore, call this view, with others, *doxastic compatibilism*. It turns out that doxastic compatibilism comes in two varieties. Some doxastic compatibilists, such as Hilary Kornblith and Matthew Chrisman, have argued that one can have doxastic obligations even if one has no control over or influence on one's beliefs whatsoever. Other doxastic compatibilists, such as Sharon Ryan and Matthias Steup, have argued that some kind of control is necessary for having doxastic obligations, but that *compatibilist* rather than *voluntary* control is enough for having such obligations.⁷ The main aim of this paper is to argue that doxastic compatibilism is untenable. I discuss the first variety of doxastic compatibilism in section 3 and the second variety in section 4. Before that, in section 2, I make a few preliminary remarks in order to get the crucial notions of responsibility and blame sharper into focus.

2. A Word on Responsibility

Let me be explicit that by 'responsibility', I do *not* mean *causal* responsibility (roughly, bringing about or contributing to the occurrence of some state of affairs), nor *ministerial* responsibility (one's being accountable for the professional actions and omissions of certain officials in one's company or department), nor *legal* responsibility (one's being the proper object of legal treatment in the form of punishment or its absence), nor *role* responsibility (we sometimes have certain tasks or duties in virtue of our role or profession),⁸ but *normative* responsibility, where I use the word 'normative' in a narrower sense than it is often used. It is hard, if not impossible, to specify this variety of responsibility in a theoretically neutral fashion—that is, in a way that does not depend on a particular account of responsibility. Maybe giving examples is the best I can do here. Florence Nightingale was normatively responsible for saving thousands of lives in the Crimean war, Adolf Eichmann was normatively responsible for the deportation of millions of Jews to concentration camps around Europe, I am normatively responsible

⁶ See, for instance, Clarke 1986; Dretske 2000; Kornblith 1983; Leon 2002; Stocker 1982. This also seems to be the view of Alston himself. In a separate section of the paper, he argues against the idea that *epistemic justification* can be understood in terms of influence on what we believe, but he does not deny that we can be *doxastically responsible* in virtue of our influence on what we believe.

⁷ It is usually this second variety of doxastic compatibilism that philosophers have in mind when they use the expression 'doxastic compatibilism'.

⁸ Thus also Hart 1970, 211-230.

for the amount of work I do on an average day and for whether I take coffee or orange juice for breakfast, a racist is normatively responsible for his heinous beliefs, and I am sometimes normatively responsible for believing certain propositions upon insufficient evidence. *Moral* responsibility falls under normative responsibility, but it is not the only kind of normative responsibility: there also seem to be such things as *prudential* responsibility and *epistemic* responsibility. In what follows, I am concerned solely with *normative* responsibility.

To be responsible, it seems to me, is to be the proper object of one or more of the *reactive* attitudes, such as praise and blame. But what are these reactive attitudes? Reactive attitudes are *affective* attitudes that we adopt primarily towards people on the basis of their actions, desires, beliefs, virtues and vices, and character. We adopt them towards other people, but also towards ourselves. Thus, we can be angry at someone's decision to remain silent and I can feel remorse about what I said to my friend last night. Reactive attitudes are as varied as blame, praise, resentment, outrage, gratitude, forgiveness, indignation, respect, compunction, and remorse. Reactive attitudes are to be distinguished from another kind of attitude that we sometimes adopt, an attitude that I call the (merely) *evaluative* attitude.⁹ Varieties of the evaluative attitude are pity, certain kinds of love, certain kinds of shame, deeming inappropriate, considering harmful, and so forth. Some evaluative attitudes are affective attitudes, whereas others are not. The crucial difference between reactive and evaluative attitudes is that in adopting the former, we hold someone responsible, whereas in adopting the latter we merely take it that the occurrence of some state of affairs was desirable, undesirable, or of neutral worth. This is not to say that if we adopt an evaluative attitude towards someone, we do *not* hold that person responsible. After all, we can adopt both a reactive and an evaluative attitude towards someone for something. It *is* to say, however, that in adopting an evaluative attitude, we do not *thereby* hold someone responsible, whereas we do if we adopt a reactive attitude.

When I say that someone is the proper object of praise or blame for φ -ing, I mean that that person *deserves* or *merits* praise or blame for φ -ing. I take this to be an objective notion: if someone deserves praise or blame, then such appraisal will be the reactive attitude that would or could be adopted by a physically, mentally, and emotionally well-functioning person who is fully informed about the situation. This means that, contrary to what some philosophers say,¹⁰ the notion of merit or desert is *not* to be spelled

⁹ I take it that Peter Strawson's *objective* attitude is identical or similar to my *evaluative* attitude (cf. Strawson 1974, 4-13). I prefer 'evaluative' to 'objective', for evaluations are often highly subjective.

¹⁰ See, for instance, Eshleman 2004.

out in terms of what a person with *justified, rational, or warranted* beliefs about the situation would do. A person with merely justified beliefs may *mistakenly* hold someone responsible, since justification and truth can come apart. It also means that a fully informed and properly functioning human being need not always actually adopt a reactive attitude towards some person *S* for something for which *S* is responsible. For instance, in considering the Nazi crimes mentioned in Simon Wiesenthal's monograph *Justice, not Vengeance*, she may be so overwhelmed by their sheer number that she will not have particular emotions towards each of them. However, being fully informed and functioning properly, she *could* adopt such reactive attitudes. Whether she will in fact adopt some reactive attitude or not depends on all sorts of situational factors.

One might wonder how we should *distinguish* reactive from evaluative attitudes. However, I do not think that providing such a criterion is necessary. In principle, we could provide an account of responsibility by including an exhaustive list of reactive attitudes. And if that list correctly picks out those attitudes that we adopt in holding people responsible, then there is no theoretical need to add a criterion to distinguish them from evaluative attitudes. For instance, it seems intuitively obvious that in resenting or blaming someone, we hold her responsible, whereas in pitying someone, we do not. And it is contradictory to say that someone is blameworthy for her bad situation but that she is *not* responsible for it, whereas it is *not* contradictory to say that someone is to pity for her bad situation but that she is *not* responsible for it.

This account is similar to the widely influential account of Peter Strawson, who also defended an account of responsibility in terms of reactive attitudes.¹¹ My account differs from that of Strawson in that I have made the additional claim that one is responsible only if one is the *proper* object of reactive attitudes and I have spelled out when someone is a proper object of reactive attitudes.

Of course, the literature also provides different accounts of responsibility. I think that, in so far as they are relevant to the issue under consideration in this paper, most of them are compatible with what I have said here.¹² However, there is one important alternative account, the *behaviourist* account, that delivers different and conflicting verdicts on certain cases that are discussed in this paper, so let me say a few words about it. On the behaviourist account, to be responsible is to be the *proper object of punishment or reward*. On this view, one is responsible to the extent that punishment and reward will increase the likelihood of good behaviour and

¹¹ See Strawson 1974.

¹² I am thinking, for instance, of *accountability* accounts (see Oshana 1997) and *ledger* accounts (see Feinberg 1970, 30-31, and Glover 1970, 64-65).

decrease the likelihood of bad behaviour in the future. This account, which has been advocated most famously by Moritz Schlick,¹³ seems to me incorrect for at least two reasons. First, we may punish and reward animals, children, and people with limited mental abilities if we believe that we can thereby train them to perform certain beneficial actions or not perform certain harmful actions. Yet, we often do not hold them responsible. Rather, we consider them as beings to be manipulated or treated in order to bring about some good or to avoid some harm. Second, there are circumstances in which we hold people responsible, but consider punishment or reward inappropriate. We hold our friends responsible for not telling what they consider to be the truth, but only in judicial circumstances are people punished for such a thing. Blame rather than punishment seems to be the appropriate reaction in such cases.

It follows from the account that I have defended that to be responsible for some belief is to be the proper object of such reactive attitudes as praise and blame for that belief. This is crucial, for we often *evaluate* people's beliefs or people for holding certain beliefs without thereby holding those people *responsible*. Consider the following sentences:

- (a) "When asked what she sees, an adult and properly functioning human being ought to believe that she is seeing a chair when she is looking straight at one."
- (b) "If Judith is to win the quiz, she should believe that Jupiter is the largest planet in our solar system."
- (c) "Since Hercule Poirot did not return to her house, he must have believed that Mrs. Burgess was not culpable."
- (d) "If she *knows* that the polar bear was reclassified as a vulnerable species in 2005, then she has to *believe* that."

In most circumstances, these will not be judgements by which we hold people responsible for some belief. Normally, (a) will be uttered to describe proper doxastic functioning, (b) to say what belief is instrumentally useful or necessary for attaining some social good, (c) to tell what belief someone can reasonably be expected to have given certain available evidence, and (d) to characterize the conceptual relation between knowledge and belief. In uttering these sentences, we need not hold the cognitive subject in question *responsible* for her belief, even if we use expressions like 'should believe', 'ought to believe', or maybe even 'has an obligation to believe'. The

¹³ See Schlick 1962, 151-158.

account of responsibility in terms of reactive attitudes that I gave above will help us to distinguish situations in which we merely *evaluate* people for holding certain beliefs from those in which we hold people *responsible* for having certain beliefs.

3. Doxastic Obligations without Doxastic Control?

On the first variety of doxastic compatibilism, Alston's argument does not count against the idea that we bear doxastic responsibility, because premise (1) is false: having doxastic obligations does not require any kind of doxastic control whatsoever. We can be praiseworthy if we meet our doxastic obligations and we can be blameworthy if we violate our doxastic obligations, whether or not we can control our beliefs. In other words, we can be responsible for our beliefs, even if we have no control over them whatsoever. Those who give this response to Alston's argument have argued that doxastic obligations or doxastic 'oughts' should be understood as role obligations, epistemic ideals, rules of criticism, or doxastic demands, and that such obligations do not require any kind of control. In this section, I argue that this response fails. There might be some sense of 'ought' or 'obligation' in which we ought to hold certain beliefs or even have an obligation to hold certain beliefs, but such oughts and obligations have nothing to do with responsibility, for in most cases they do not seem to imply blameworthiness in case one does not meet such oughts or obligations. To the extent that they do, they *do* seem to require some kind of doxastic control and, hence, are not convincing counter-examples to premise (1) of Alston's argument.

3.1. Doxastic Obligations as Role Obligations

First, according to Richard Feldman, doxastic obligations are *role obligations*. Parents ought to take care of their children, teachers ought to explain things clearly, and cyclists ought to cycle well, whether they are able to do so or not. Similarly, judgements on belief prescribe the right way to play the role of a believer, even if one has no control over one's belief.¹⁴ Feldman confines his account to *epistemic* doxastic obligations, that is, "evaluations that have more to do with epistemologically central matters such as knowledge and rationality".¹⁵ The right way to play the role of a believer, according to Feldman, is to believe in accordance with one's evidence. One might object that there is a crucial difference between the role of believer on the one hand and the roles of, say, teacher and cyclist on the other. For, in opposition to the former role, one has the latter roles

¹⁴ See Feldman 2000, 676-677.

¹⁵ Feldman 2008, 346.

voluntarily. Feldman agrees, but points out that there are many other roles that are as involuntary as that of believer, such as the roles of eater and breather. That these roles are involuntary does not imply that there are no correct ways to eat or breathe.¹⁶

I think that, unfortunately, the analogy fails. We *do* have control over the way we eat or breathe. To the extent that we do not, we are *not* the proper object of praise or blame for eating or breathing. It seems unfair to hold me responsible for the way I breathe if there is nothing I can do to change it. Perhaps we can hold a teacher responsible for teaching badly, even if he could not do any better. If so, however, that would be because the role of teacher is a role he has *voluntarily* accepted. Thus, to the extent that role obligations imply responsibility, they require some kind of control. Such control is remarkably absent in the case of being a believer: I have control neither over my beliefs nor over my being a believer.

Feldman points out that we sometimes praise a person for things that are not under her control, such as her beauty.¹⁷ I agree, but we should distinguish between praise as a *reactive* attitude from praise as an *evaluative* attitude. I might praise Miranda for her beauty or my recently bought Chevrolet for its speed, but I do not thereby hold them responsible for these qualities. We do not hold Miranda responsible for her beauty, in the same way as we do not blame someone else for being ugly, unless, of course, they had control over being beautiful or being ugly (but then, such praise would not count against the idea that responsibility requires control). Such merely evaluative praise ought to be clearly distinguished from praise as a *reactive* attitude. If we adopt the *reactive* attitude of praise towards someone for ϕ -ing, we can *in principle* also, say, blame or resent that person for ϕ -ing.

3.2. Doxastic Obligations as Epistemic Ideals

Second, Hilary Kornblith construes doxastic obligations in terms of ideals that take into account human limitations. More specifically, he says that what he has in mind are *epistemic* ideals. Epistemic ideals take into account what humans can believe, but they are not confined to what a particular human can believe: sometimes one ought to believe a proposition that one cannot believe. Epistemic ideals can be reached by at least *some* human beings. Doxastic obligations, therefore, provide the middle ground between epistemic ideals that are insensitive to human capacities and ideals that are

¹⁶ See Feldman 2008, 351. Feldman 2000, 674, distances himself from his earlier view on which doxastic obligations are contractual obligations (see Feldman 1988, 240-243).

¹⁷ See Feldman 2000, 676.

so constrained by a particular individual's capacities that they are unworthy of pursuit. Since being subject to epistemic ideals does not require control, this approach could save doxastic obligations.¹⁸

It seems to me that this account is flawed. The problem with our beliefs is that they seem under *no one's* voluntary control. Thus, even if we confine ourselves to epistemic ideals that can be reached only by those who do epistemically well, there will not be any doxastic obligations. No matter what the ideal attitude is towards a proposition, if there is no way I or anyone else can voluntarily acquire that attitude, it seems clear that I have no obligation to acquire that attitude. Again, we may very well be the proper object of *evaluative* attitudes for not reaching certain epistemic ideals. It follows neither that we are the proper object of *reactive attitudes* nor that we have doxastic obligations. Of course, one could say that we have an obligation to believe that p if and only if one ought to believe that p and that one ought to believe that p if and only if the epistemic ideal is to believe that p . Perhaps the word 'obligation' is sometimes used in this contrived sense. But then, it follows from what I argued in the previous section that we will no longer be talking about doxastic *responsibility*.

3.3. Doxastic Obligations as Rules of Criticism

Third, Matthew Chrisman argues that we should interpret doxastic obligations as what Wilfrid Sellars calls *rules of criticism* that materially imply *rules of action*. Whereas rules of criticism (ought-to-be's) concern ways of being, rules of action (ought-to-do's) concern actions. One can be subject to a rule of criticism for ϕ -ing, even if one has no control over ϕ -ing, whereas rules of action *do* require control. The material implication which Chrisman has in mind can be spelled out as follows:

- (4) If X ought to be in state Φ , then, other things being equal and where possible, one ought to bring it about that X is in state Φ .

Now, as Chrisman rightly notices, (4)'s consequent can be interpreted in three different ways: (i) on the *conditional view*, X herself ought to do what she can to bring it about that X is in Φ , (ii) on the *universal view*, everyone ought to do what she can to bring it about that X is in Φ , and (iii) on the *existential view*, someone ought to do what she can to bring it about that X

¹⁸ See Kornblith 2001, 238-239. In a previous article (Kornblith 1983, 33), he distinguishes between doxastic obligations as epistemic ideals and doxastic obligations that imply responsibility. In an even earlier paper, he understands doxastic justification in terms of the absence of epistemic culpability (Kornblith 1982, 243). Unfortunately, the distinction between ideals and responsibility is absent from Kornblith 2001 and 2002, 137-161.

is in Φ . According to Chrisman, which view is correct depends on which kind of ought is involved. For instance, the rules of criticism

(e) people ought to feel outrage about genocide;

(f) this child ought to be able to tie his shoes by age four;

materially imply respectively a universal and an existential rule of action:

(e') everyone ought to do what she can to bring it about that people feel outrage about genocide;

(f') someone ought to do what she can to bring it about that this child is able to tie his shoes by age four.¹⁹

According to Chrisman, the *epistemic* ideal for humans is to be good information tracking and transmitting beings. Ought-to-do's implied by doxastic ought-to-be's are interpersonal and sometimes intrapersonal forward-looking and backward-looking rules of action on the part of one's epistemic community, sometimes including oneself. Thus,

(g) you ought to disbelieve that the earth is flat;

is an ought-to-be that materially implies the following correlative existential ought-to-do:

(g') your parents and teachers ought to have taught you that the earth is not flat.²⁰

If this is correct, doxastic obligations do not require control on the cognitive subject's part. It seems to me, however, that this view is problematic. Chrisman rightly acknowledges that there are many situations in which *S* in some sense ought to believe that *p*, without *S*'s herself being subject to a materially implied ought-to-do. But then it is incorrect to say that *S herself* is subject to an *obligation to believe that p* or that she is *blameworthy* if she fails to meet that obligation. Assuming that there are such things as epistemic obligations on the side of one's epistemic community, one's epistemic community may have an obligation to convince one of the truth of *p* and may be blameworthy for failing to do so, but, clearly, that is something different. The sense in which *S* ought to believe that *p* in such situations, then, will *not* be normative, but merely evaluative.

¹⁹ See Chrisman 2008, 358-363.

²⁰ See Chrisman 2008, 369-370.

3.4. Doxastic Obligations as Doxastic Demands

Finally, according to Philippe Chuard and Nicholas Southwood, in expressing that we hold someone responsible for a belief, we make judgements that are relevantly like judgements such as:

- (h) Oscar ought to feel guilty for what he did to his sister, and
- (i) Judy ought to understand what Nicole is going through,

in that these judgements make *demands* on us. They demand that we respond in certain ways. But making a demand to φ on some person S does not presuppose that S has control over her φ -ing, only that S can φ . The idea that S can φ can be spelled out in terms of there being the logical possibility that S φ -s, in terms of S 's having the alternate possibilities of φ -ing and $\sim\varphi$ -ing, or in terms of S 's having the capacity to φ . But none of these implies that S has *voluntary control* over φ -ing. Alston's argument goes wrong, then, in conflating the idea that S can φ with the idea that S has control over φ -ing. Hence, doxastic obligations do not require control.²¹

It seems to me that this approach faces a serious difficulty. For it is not clear what it is to be the proper object of a *demand*. If it does *not* have to do with being responsible, then Chuard's and Southwood's strategy does not even address the argument from lack of doxastic control that William Alston and others give. If it *has* to do with being responsible, then it follows from what I argued in section 2 that if a demand is made on someone, then that person is blameworthy if she fails to meet the demand (if she has no good excuse). Chuard and Southwood, however, explicitly reject the idea that doxastic obligations are in any way relevantly related to blameworthiness. They do so, because they do not find any plausible interpretation of blame on which blameworthy belief requires doxastic control. As they see it, blame can be interpreted in terms of being criticisable and in terms of other people's having certain legitimate expectations. As they rightly point out, being criticisable does not require control. And someone else's having a legitimate expectation about what one will believe is an implausible way of spelling out doxastic blame, for usually people have no expectations about what we will believe.²²

However, they overlook the option that I defended in section 2, namely that one is blameworthy if and only if one is the proper object of a negative reactive attitude, such as resentment or blame. I think that sentences (h) and (i) can plausibly be understood along these lines, that is, as expressing reactive attitudes towards Oscar and Judy. But if they do, then it seems that a

²¹ See Chuard and Southwood 2009, 601, 614-619.

²² See Chuard and Southwood 2009, 620-623.

rational speaker would assume that there is something Oscar and Judy could have done about respectively not feeling guilty and lacking understanding. If there is nothing they could possibly do or have done to feel guilty or to understand, then it is hard to see how we could properly blame them. How, for instance, can we rightly blame Oscar for not feeling guilty for what he did to his sister if he has tried everything he possibly could, but still fails to feel guilty? Thus, a person is properly subject to some demand only if that person has control over meeting that demand. However, as we saw, there is good reason to think that we have no voluntary control over our beliefs.

3.5. *Intermediate Conclusion*

In this section, I have criticized four attempts to argue that doxastic obligations belong to a special kind of obligations that do not require control. For each of the four kinds of obligations that I discussed—role obligations, epistemic ideals, rules of criticism, and doxastic demands—I agree that there may be some sense of the word ‘obligation’ in which we have doxastic obligations along these lines or at least that it can be true that we ought to hold some belief. In all four cases, however, it turned out that such *oughts* have nothing to do with responsibility. And that is because they have nothing to do with praise, blame, or any of the other reactive attitudes. To the extent that they do have to do with these attitudes, they seem to require some kind of control over or influence on our beliefs. Of course, this does not establish that there is no way that one could plausibly argue that doxastic obligations are of a special kind that does not require control. But, as far as I know, these four proposals are all the attempts that we find in the literature to argue that doxastic obligations do not require any kind of doxastic control. I take it, though, that my criticisms of these four proposals provide sufficient reason to seek an alternative solution to the problem formulated by William Alston.

4. **Doxastic Obligations in Virtue of Doxastic Compatibilist Control?**

4.1. *Compatibilist Doxastic Control*

The second variety of doxastic compatibilism also denies premise (1), but for a different reason. The idea is *not* that doxastic obligations do not require any kind of control whatsoever, but that doxastic obligations do not require *voluntary* control, as I defined ‘voluntary control’ earlier. Remember that someone has voluntary control over φ -ing if and only if one can φ as the result of an intention to φ and one can $\sim\varphi$ as the result of an intention to $\sim\varphi$. Most doxastic compatibilists of this second variety grant that we cannot form our beliefs as the result of an intention to do so, although, as

we shall see, Matthias Steup is an exception. All of them agree, though, that it is *not* required that *in exactly the same evidential circumstances*, one could have chosen to hold a different belief.

What does doxastic compatibilist control amount to? Doxastic compatibilists answer this question differently, although their accounts show important similarities. Let me briefly characterize the three main accounts that we find in the literature. First, according to Mark Heller, people have an *epistemic nature*, that is, second-order desires to form beliefs in accordance with certain dispositions rather than others. What is required for doxastic control is that our beliefs reflect our epistemic nature, that is, that we form the beliefs we form because of the epistemic nature we have. We are responsible in such cases because our beliefs manifest who we are from an epistemic point of view.²³

Second, according to Sharon Ryan, what is required for doxastic responsibility is that we can appreciate evidence and form beliefs in accordance with it, in the same way as we can weigh various practical considerations and act in accordance with them. We are responsible for our beliefs if they are unlike typical coerced actions and like actions such as typing the letters that we type and moving our limbs when we have been running for a while. These are actions that are responsive to reasons, but do not seem to involve the formation of any intentions. It does not require an explicit intention to perform them. It follows that intentional doxastic control is *not* required for having a doxastic obligation.²⁴

Third, according to Matthias Steup, some person is responsible for a belief if that belief is the outcome of a process that is responsive to epistemic reasons, i.e. evidence, and if that person's belief is *weakly intentional*. Something is weakly intentional if it is non-accidental and if one has a pro-attitude towards it. Thus, my stepping on the clutch is weakly intentional if it does not result from such things as a sudden cramp in my leg and if I mean to step on the clutch. A belief is weakly intentional if it is not due to cognitive malfunction and if one endorses it or if one is comfortable with it. Beliefs, then, are like actions such as starting a car by inserting the ignition key, engaging the clutch, shifting into reverse, and stepping on the gas. These automatic, unthinking, and habitual actions are not performed as the result of an explicit intention, but are nevertheless under one's control.²⁵

²³ See Heller 2000, 132-137.

²⁴ See Ryan 2003, 70-74. An account similar to that of Ryan is Owens 2000, 115-129. His account is different in that he is not willing to describe reasons-responsiveness as control; he simply denies that doxastic responsibility requires any kind of control.

²⁵ See Steup 2012. For other, less developed compatibilist accounts of doxastic control in the same spirit, see Hieronymi 2008, 362-363; Jäger 2004, 217-227; Shah 2002, 443; Smith 2000, 40-46; Smith 2005, 236-271; Weatherson 2008, 546.

One may object that these are not really instances of doxastic *control*. But the doxastic compatibilist will simply disagree and will affirm that this is all that is required for doxastic control. My criticism on this variety of doxastic compatibilism that I provide in the following section will, therefore, not depend on whether or not compatibilist doxastic control can properly be called ‘control’. Rather, I will argue that compatibilist doxastic control is insufficient for being responsible for one’s beliefs.

4.2. A Problem with Doxastic Compatibilist Control

I think the prospects for doxastic compatibilism of this second variety are as bleak as those for the first variety. If, on the one hand, one *has* some kind of control (whether intentional or otherwise) over or influence on one’s higher-order beliefs or one’s reasons-responsive processes, then it seems that in the scenarios sketched by doxastic compatibilists, one is responsible for one’s beliefs *not* because one has some kind of doxastic compatibilist control over them, but because one *influences* what one believes by having *control* over such belief-influencing factors as higher-order beliefs and reasons-responsive processes. If, on the other hand, one *lacks* some kind of control over or influence on one’s higher-order beliefs and reasons-responsive processes, then how can one be responsible for them and for the ensuing beliefs?

It seems to me that most, if not all doxastic compatibilists are aware that the first horn of this dilemma is problematic and, therefore, opt for the second horn. This horn, however, is equally problematic. Imagine that Nagoni is raised in a culturally isolated community. From an early age onward, she is indoctrinated in the tradition of the tribe. She, therefore, believes that she should believe anything that the tradition teaches and so she desires to believe anything that is part of her tradition. Thus, upon considering the proposition that humans are fallen angels, she immediately believes this proposition, for this is one of the core teachings of her tribe. She believes this as strongly as anything she could possibly believe. It is completely irresistible. Since she is convinced that this is what the tradition teaches, she believes in accordance with her epistemic nature. But it is clear that she has no control over her belief: she has been indoctrinated to such an extent that her belief is genuinely irresistible. This shows that Heller’s account in terms of epistemic natures is untenable. The reasons-responsiveness accounts of Ryan and Steup do not face this problem, for Nagoni’s belief-forming mechanism clearly is not reasons-responsive: she would hold that belief no matter what her evidence were.

Ryan’s and Steup’s accounts face another problem, though. Imagine a possible world in which there are creatures who are like us in that their belief-forming mechanisms are largely functioning properly: upon having

the experiences and beliefs we have, they roughly form the same beliefs as we do. Moreover, they are as responsive to evidential reasons as we are: upon having (significantly) different evidence, they form different beliefs. Once their evidence changes, their beliefs change accordingly. In one regard, however, they are crucially different from us: they cannot *influence* what they believe. Thus, they cannot gather evidence, work on their intellectual virtues and vices, improve the functioning of their cognitive mechanisms, and so forth. For instance, they cannot decide to gather further evidence on a politically sensitive issue or decide to try to become more open-minded. Nonetheless, their beliefs change every now and then, because the doxastic mechanisms producing those beliefs are responsive to changes in the evidential situation. In fact, they are caused in a paradigmatically good way: perceiving a tree causally leads to the belief that there is a tree in front of them, tasting something sweet causally leads to the belief that they are tasting something sweet and when that taste changes to something bitter, they come to believe that they are tasting something bitter.

Would we hold those creatures responsible for their beliefs in this scenario? It seems clear to me that we would not. Their belief-formation is clearly not up to them. Their beliefs are simply the deliverances of their cognitive mechanisms in combination with certain inputs. But, we have assumed, neither the functioning of their cognitive mechanisms nor the scope or quality of their evidence base is up to them. It seems clear that if these are not up to them, the output is not up to them either and it would be unfair to hold them responsible for their beliefs. I submit, then, that the plausibility of Ryan's and Steup's doxastic compatibilism derives from the fact that in the scenarios they sketch, people have control over all sorts of factors that *influence* what the subjects believe. It is because we can *intentionally* perform such belief-influencing actions and *not* because our belief-forming mechanisms respond differently to different inputs (what they call 'compatibilist doxastic control') that we can be held responsible for our beliefs in such cases. But, as I pointed out in section 1, that would count as a *different* response to Alston's argument, a response that denies that it follows from (3)—the thesis that we do not have doxastic obligations—that we are not responsible for our beliefs. We would *not* be responsible in virtue of our presumed doxastic compatibilist control over our beliefs.

4.3. Three Arguments for Doxastic Compatibilism and Three Replies

The argument that I just gave against the second variety of doxastic compatibilism counts against any view on which we are responsible for our beliefs in virtue of our compatibilist control over our beliefs. For, it seems that if someone is merely reason-responsive without having the ability to *influence* her beliefs, she is not at all responsible for her beliefs. Hence, compatibilist

control can never suffice for doxastic responsibility. However, we should hold on to the idea that we cannot have obligations to believe particular propositions unless we have sufficient voluntary control (as I defined it earlier on) over our beliefs *only if* such a view does not face equally insurmountable difficulties. Let me, therefore, discuss the three main arguments that compatibilists have levelled against the idea that doxastic obligations require doxastic *voluntary* control rather than doxastic *compatibilist* control, where ‘voluntary control’ should be understood as I defined it earlier, namely implying that one has alternate possibilities.

First, one might think that what counts both when it comes to action and when it comes to belief is *reasons* rather than *intentions*. Action and belief are similar in that if one takes oneself to have sufficient reason to perform act *A* or hold belief *B*, one will normally perform *A* or hold *B*. For instance, if I have convincing reasons to slow down and stop when I approach an intersection, I will stop, and if I have convincing reasons not to stick a knife in my arm, I will not do so. Of course, I could decide to act differently if I had an overriding reason to do so, but similarly I could *believe* differently if my epistemic reasons were different.²⁶ I find this reply unconvincing. First, if I take myself to have good reasons to do *A*, I will not do *A* as long as I do not also *intend* to do *A* (unless doing *A* is a habitual action; I return to that momentarily). Reasons alone, then, will not suffice to explain action: we also need intentions or, as I shall explain below, at least the *ability* to do the thing in question intentionally.²⁷ Second, there are many situations in which my practical reasons for doing *A* and *not* doing *A* are balanced. In such situations I can deliberate about whether to do *A* or *not* to do *A* and I can consequently do *A* or *not* do *A*. For instance, I can choose to put cheese or peanut butter on my bread.²⁸ However, in all or most situations in which my evidence for and against *p* is balanced, I cannot simply decide to believe or disbelieve that *p*. Rather, I automatically find myself with a particular doxastic attitude towards *p*, normally that of withholding (withholding both belief and disbelief).

Second, one could object that we are responsible for and have control over actions such as my typing the letters that I type, my moving my legs

²⁶ For this argument, see Russell 2001, 42-43; Ryan 2003, 63-64; Steup 2000, 46, 54; 2001, 17n.

²⁷ According to Steup 2012, we *can* form our beliefs voluntarily. I return to this thesis below.

²⁸ Of course, if determinism is true, it is *not* the case that there are two options available to one. Even if determinism is true, though, it will seem to the cognitive subject in question that it is up to her whether she does *A* or she does *not* do *A*. When one’s evidence regarding *p* is balanced, however, it will seem to one that one *cannot* believe that *p* and that one *cannot* believe that $\sim p$; it will seem to one that, given one’s evidence, one cannot but suspend judgement on *p*.

when I have been running for a while, and my stepping on the clutch, even though we often do not perform these actions as the result of a preceding intention to do so. Hence, intentions are not crucial to control and responsibility.²⁹ In response, let me point out that beliefs are crucially different from such actions as moving my legs and typing the letters I type. I can type the letters on my computer without explicitly intending to do so, only because I once did intend to type these letters, and something similar applies to actions like running. I do not think that this is true for *all* actions that are under our control. I can exercise control over breathing or not breathing, even if I have never *intentionally* done so. But breathing or not breathing—or, at least, breathing or not breathing at some particular moment—is something that I could *in principle* do intentionally. With belief, things are different. I do not now unintentionally form certain beliefs because I once intentionally formed them. Nor is it the case that I could *in principle* form a belief as the result of an intention to do so. Rather, I have always found myself with certain beliefs, given the evidence I had. Beliefs may be quite different from compelled actions, but they are not very much like the unintentional actions just mentioned either.

Third, one might object that the idea that control requires the ability to form intentions becomes problematic if we apply this idea to control over intentions themselves. Two problems have been identified here. First, we lack intentional control over our intentions. We cannot form intentions to φ as a result of deciding to intend to φ upon believing that intending to φ would be good. We can intend to φ only if we believe that φ -ing itself is good. We are deluded into thinking that we have intentional control over our intentions, because an intention is usually good to have just in case the action intended is good to perform, but there are cases in which the goodness of intentions and that of the intended actions come apart.³⁰ Second, the idea that control requires the ability to form intentions leads to an infinite regress. One's intention to φ is under one's control only if one could form an intention to intend to φ . But that intention would be under one's control only if it could be formed by a further intention, and so on. In order to avoid this infinite regress, we should deny that control over φ -ing requires the ability to φ intentionally.³¹ And if these two points about intentions are correct, then why would we not think that beliefs are like intentions in that they are under our control but not in virtue of a presumed ability to form an intention to have them?

Let me first address the final claim, that is, the assertion that beliefs are like intentions. I think that this claim is false. If I have equally good reason

²⁹ For this objection, see Steup 2008, 284-285; 2011, 154-155.

³⁰ This point has been made by Owens 2000, 81; Hieronymi 2006, 56-57; 2008, 368-371.

³¹ For this line of reasoning, see Shah 2002, 440-442.

to do *A* as *not* to do *A* and equally good reason to intend to do *A* as to intend *not* to do *A*, as is often the case, then I can equally well intend to do *A* as intend *not* to do *A*. For example, if I do not mind whether I take coffee or orange juice for breakfast, I can equally well intend to take coffee as I can intend to take orange juice. Whether I intend to do some action *A* or *not* to do *A* is in such cases up to me, in a way that it is not up to me what doxastic attitude I take towards *p* if my evidence bearing on *p* is balanced. Thus, it seems that I have a kind of control over my intentions that I lack over my beliefs, however precisely this kind of control is to be spelled out. As to the two points about intentions, I agree that one's control over one's intentions is not to be spelled out in terms of one's ability to form intentions to intend. I submit that some person *S*'s intention is under *S*'s control simply if it is formed by *S* herself rather than anything else, such as an evil neurosurgeon or a brain tumour. To form such an uncoerced intention is to exercise one's will. One might wonder why, if this is true, we should not say that, similarly, to form a belief is to exercise one's will. I answer that to will something simply *is* to intend to do it, to choose to do it, to decide to do it. To *believe* something is *not* to will something. Intentions seem to be essential to exercising one's will in a way that one's beliefs seem not to be.

One may reply, as Steup does in his 2008 paper, by asking why we should give pride of place to intentions. Actions are typically caused by intentions and desires, whereas beliefs are typically caused by our evidence. How is it supposed to follow that beliefs are not under our control? The two domains of actions and beliefs have different criteria of control.³² To say that control over φ -ing requires the ability to form an intention to φ is to measure the two realms with the same yardstick, whereas they should be measured with different yardsticks. In response, I would like to stress that the intuition that I share with many philosophers that control over φ -ing requires the ability to intend to φ is not an intuition about actions, but about the very nature of control: one cannot have control over φ -ing if one is unable to form an intention to φ .³³ It seems implausible that control in one realm is entirely different from control in another realm if, as Steup and Ryan agree, to be responsible is to be the proper object of reactive attitudes like praise and blame, and if we bear responsibility in virtue of having control. It seems that we should reject such a view unless we find a plausible candidate of control in the doxastic realm that seems to render us doxastically responsible. As I argued above, doxastic compatibilists of the second variety have not put forward any viable candidate.

Another option, which Steup defends in a more recent paper (Steup 2012), is to claim that we *do* form beliefs as the result of an intention to do

³² See Steup 2008, 387-390.

³³ Thus also Booth 2009, 9-11.

so. Thus, we form beliefs in a way that is *not* significantly different from how we perform actions as the result of an intention to do so. According to Steup, one can, in believing that *p*, carry out an intention to believe that *p*. Consider the following case that Steup sketches:

having returned from a trip and taken a shuttle to the airport parking garage, I am now where I thought I left my car. To my surprise, it is no longer there. I wonder whether it has been stolen. There is of course the possibility that I don't accurately remember where I parked it. So I retrieve the paper slip which states the exact location of my parking spot. According to the slip, I am at the right spot. Considering my evidence—the parking slip and the absence of my car—I conclude that it was stolen.³⁴

According to Steup, a situation like this is plausibly described by saying that I come to believe that my car was stolen because, considering my evidence, I *decided* to believe that my car was stolen (and the causal relation between this belief and the decision to hold it is non-deviant). The decision to believe that my car was stolen is, according to Steup, analogous to deciding to take a walk: one considers one's reasons for (not) taking a walk and then decides to take a walk. Steup discusses and replies to two arguments against this interpretation of the car theft scenario. Here, I will levy two objections of my own against Steup's interpretation. For, there are at least two important differences between coming to believe that my car was stolen and deciding to take a walk.

First, Steup is right that in both scenarios, before considering my reasons, it seems to me that there are two options; namely, respectively, believing that *p* and *not* believing that *p*, and taking a walk and *not* taking a walk. However, the scenarios are different *once I have considered my reasons*. When I have considered my reasons for (not) believing that the car was stolen and find myself with sufficient evidence for believing that it was, I cannot but believe that the car was stolen. But when I have considered my reasons for (not) taking a walk and find myself with equally good reasons to take a walk as to *not* take a walk—say, because the weather is really nice, but I also do not have that much time left to prepare a lecture that I should deliver this evening—it seems it is *up to me* whether or not I take a walk. This is confirmed by how we interpret such scenarios afterwards. We think that, given the exact same practical reasons, we could have decided *not* to take a walk. But we do *not* think that, given the exact same evidential reasons, we could have *failed* to believe that the car was stolen (or, at least, we would deem the absence of such a belief irrational).

³⁴ Steup 2012, 157.

We should notice that this remains true when we sketch two scenarios that are even more analogous by coming up with a situation in which I have to decide what to believe, but in which my evidential reasons for and against p are balanced (in the same way as my reasons for and against taking a walk are balanced). Here, one could think of our evidence regarding such propositions as that the number of stars is even or that the next time I flip this coin, it will turn up heads. *Once I have considered all my reasons*, I cannot but suspend judgement on the proposition that the number of stars is even and on the proposition that the coin will turn up heads next time I flip it. But, again, once I have considered all my reasons for and against taking a walk, it seems that it is still up to me whether or not I take a walk.

Second, it is an important characteristic of voluntary actions that we can decide *when* we perform the action in question. I can now decide to take a walk in ten minutes, or in half an hour, or tomorrow. In deciding whether or not to perform an action, I also decide *when* I will perform it or, at least, whether or not I will perform it right now (I might decide to perform it *at some time* in the future, without deciding on exactly *when* I am going to perform it). When I decide to go to the library, to buy a car, or to raise my arm, I thereby also decide whether to do it straight away or to do it later and if the latter, we often decide *when* we will do it. If I decide to purchase a car tomorrow, then from that moment onwards—unless I change my mind—I explicitly or implicitly *intend* to purchase a car tomorrow. All this is absent in the case of belief. Once I have considered my evidence regarding p , I find myself automatically and immediately believing that p , disbelieving that p , or suspending judgement on p . I cannot decide to believe that p later—in, say, ten minutes or tomorrow. Nor can I have the intention for an extended period of time to believe that p (although I may well have the intention to *try to acquire the belief that p* during a certain period of time). This is a second, important difference between the process of coming to believe that my car was stolen and deciding to take a walk. It seems to me that, jointly, these two differences provide us with enough reason to say that, even though we decide to *consider our evidence* regarding certain propositions and, therefore, decide to *weigh our epistemic reasons*, we do *not* decide to form a belief or come to hold a belief as the (non-deviant) result of an intention to do so.³⁵

5. Conclusion

In this paper, I have discussed doxastic compatibilism, which is one of the major responses (or maybe *the* major response) that has been given to

³⁵ In this section, I have argued that doxastic compatibilist control does not suffice for doxastic obligations. I have *not* provided a defence of the thesis that doxastic obligations require the ability to believe otherwise. For a defence of that thesis, see Peels 2013.

Alston's well-known argument against the deontological conception of epistemic justification but which also counts against the thesis that we are at least sometimes responsible for our beliefs. Doxastic compatibilists claim that we can have doxastic obligations, even if we have no voluntary control over our beliefs, and that we can be praiseworthy and blameworthy for our beliefs in virtue of meeting and violating those obligations. According to some doxastic compatibilists, having doxastic obligations requires no doxastic control whatsoever. According to other doxastic compatibilists, having doxastic obligations does not require *voluntary* doxastic control (the control one has when one can choose to do *A* and one can choose *not* do *A*), but only *compatibilist* doxastic control (roughly, reason-responsiveness). I have argued that the former confuse doxastic *responsibility* with closely related phenomena, whereas the latter use examples which derive their plausibility from a view which turns out to be quite different from doxastic compatibilism, namely that we are responsible for our beliefs in virtue of the fact that we can control factors that make a difference to what we believe, such as our evidence base. If my arguments are convincing, then we have good reason to think that doxastic compatibilism is untenable and that, in order to save the thesis that we are doxastically responsible, we should argue that we can voluntarily choose our beliefs after all or that we bear doxastic responsibility in virtue of our control over factors that make a difference to what we believe, such as our evidence base and our intellectual virtues and vices, rather than over those beliefs themselves.³⁶

References

- Alston, William P. (1989). "Concepts of Epistemic Justification", in *Epistemic Justification: Essays in the Theory of Knowledge* (Ithaca: Cornell University Press), 81–114.
- (2005). *Beyond "Justification": Dimensions of Epistemic Evaluation* (Ithaca: Cornell University Press).
- Booth Anthony, R. (2009). "Compatibilism and Free Belief", *Philosophical Papers* 38.1, 1–12.
- Chrisman, Matthew. (2008). "Ought to Believe", *The Journal of Philosophy* 105.7, 346–370.
- Chuard, Philippe and Nicholas Southwood. (2009). "Epistemic Norms without Voluntary Control", *Noûs* 43.4, 599–632.

³⁶ For their helpful comments on presentations or earlier written versions of (sections of) this paper, I would like to thank Anthony Booth, Jeroen de Ridder, Richard Feldman, Hilary Kornblith, Conor McHugh, Philip Nickel, Herman Philipse, Michael Smith, Nicholas Southwood, Ralph Wedgwood, René van Woudenberg, Nicholas Wolterstorff, and an anonymous referee of this journal.

- Clarke, Murray. (1986). "Doxastic Voluntarism and Forced Belief", *Philosophical Studies* 50.1, 39–51.
- Dretske, Fred. (2000). "Entitlement: Epistemic Rights without Epistemic Duties?", *Philosophy and Phenomenological Research* 60.3, 591–606.
- Eshleman, Andrew. (2004). "Moral Responsibility", *The Stanford Encyclopedia of Philosophy*, <http://plato.stanford.edu/>, first published January 6th 2001, substantive revision November 18th 2009, last visited: November 7th 2012.
- Feinberg, Joel. (1970). *Doing and Deserving: Essays in the Theory of Responsibility* (Princeton: Princeton University Press).
- Feldman, Richard. (1988). "Epistemic Obligations", in James E. Tomberlin (ed.), *Philosophical Perspectives 2: Epistemology* (Atascadero, CA: Ridgeview Publishers), 235–256.
- . (2000). "The Ethics of Belief", *Philosophy and Phenomenological Research* 60.3, 667–695.
- . (2008). "Modest Deontology in Epistemology", *Synthese* 161.3, 339–355.
- Glover, Jonathan. (1970). *Responsibility* (London: Routledge and Kegan Paul).
- Hart, Herbert L.A. (1970). *Punishment and Responsibility: Essays in the Philosophy of Law* (Oxford: Clarendon Press).
- Heil, John. (1992). "Believing Reasonably", *Noûs* 26.1, 47–61.
- Heller, Mark. (2000). "Hobartian Voluntarism: Grounding a Deontological Conception of Epistemic Justification", *Pacific Philosophical Quarterly* 81.2, 130–141.
- Hieronymi, Pamela. (2006). "Controlling Attitudes", *Pacific Philosophical Quarterly* 87.1, 45–74.
- . (2008). "Responsibility for Believing", *Synthese* 161.3, 357–373.
- Jäger, Christoph. (2004). "Epistemic Deontology, Doxastic Voluntarism, and the Principle of Alternate Possibilities", in Winfried Löffler and Paul Weingartner (eds.), *Knowledge and Belief: Proceedings of the 26th International Wittgenstein Symposium* (Vienna: Öbv&hpt Verlagsgesellschaft), 217–227.
- Kornblith, Hilary. (1982). "The Psychological Turn", *Australasian Journal of Philosophy* 60.3, 238–253.
- . (1983). "Justified Belief and Epistemically Responsible Action", *Philosophical Review* 92.1, 33–48.
- . (2001). "Epistemic Obligation and the Possibility of Internalism", in Abrol Fairweather and Linda Zagzebski (eds.), *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility* (Oxford: Oxford University Press), 231–248.
- . (2002). *Knowledge and Its Place in Nature* (Oxford: Clarendon Press).

- Leon, Mark. (2002). "Responsible Believers", *The Monist* 85.3, 421–435.
- Nottelmann, Nikolaj. (2007). *Blameworthy Belief: A Study in Epistemic Deontologism* (Dordrecht: Springer).
- Oshana, Marina A.L. (1997). "Ascriptions of Responsibility", *American Philosophical Quarterly* 34.1, 71–83.
- Owens, David. (2000). *Reason without Freedom: The Problem of Epistemic Normativity* (London: Routledge and Kegan Paul).
- Peels, Rik. (2013). "Does Doxastic Responsibility Entail the Ability to Believe Otherwise?", forthcoming in *Synthese*.
- Price, Henry H. (1954). "Belief and the Will", *Proceedings of the Aristotelian Society*, Supplementary Volume 28, 1–26.
- Russell, Bruce. (2001). "Epistemic and Moral Duty", in Matthias Steup (ed.), *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue* (Oxford: Oxford University Press), 34–48.
- Ryan, Sharon. (2003). "Doxastic Compatibilism and the Ethics of Belief", *Philosophical Studies* 114.1–2, 47–79.
- Schlick, Moritz. (1962). *Problems of Ethics* (New York: Dover Publications).
- Shah, Nishi. (2002). "Clearing Space for Doxastic Voluntarism", *The Monist* 85.3, 436–445.
- Smith Angela, M. (2000). "Identification and Responsibility", in Ton van den Beld (ed.), *Moral Responsibility and Ontology* (Dordrecht: Kluwer Academic Publishers), 233–246.
- . (2005). "Responsibility for Attitudes: Activity and Passivity in Mental Life", *Ethics* 115.2, 236–271.
- Steup, Matthias. (2000). "Doxastic Voluntarism and Epistemic Deontology", *Acta Analytica* 15.1, 25–56.
- . (2001). "Introduction", in Matthias Steup (ed.), *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue* (Oxford: Oxford University Press), 3–18.
- . (2008). "Doxastic Freedom", *Synthese* 161.3, 375–392.
- . (2012). "Belief Control and Intentionality", *Synthese* 188.2, 145–163.
- Stocker, Michael. (1982). "Responsibility Especially for Beliefs", *Mind* 91.363, 398–417.
- Strawson Peter, F. (1974). "Freedom and Resentment", in *Freedom and Resentment and Other Essays* (London: Methuen), 1–25.
- Van Woudenberg, René. (2009). "Responsible Belief and Our Social Institutions", *Philosophy* 84.327, 47–73.
- Weatherson, Brian. (2008). "Deontology and Descartes' Demon", *The Journal of Philosophy* 105.9, 540–569.
- Wolterstorff, Nicholas. (2010). *Practices of Belief*, Selected Essays, Vol. 2, ed. Terence Cuneo (Cambridge: Cambridge University Press).