

WHY RESPONSIBLE BELIEF IS PERMISSIBLE BELIEF*

RIK PEELS

VU University Amsterdam

ANTHONY BOOTH

*University of Sussex***1. Introduction**

It is widely thought that we are at least sometimes responsible for our doxastic attitudes. We can be praised, blamed, or neutrally appraised for our beliefs, disbeliefs, and withholdings. However, many epistemologists have also pointed out that, despite the fact that we are responsible both for certain actions and for certain beliefs, there are crucial differences between responsibility for beliefs and responsibility for actions. In this paper we discuss one such alleged difference, namely that whereas in many circumstances two or more different actions are *equally permissible*, our evidence is always such that there is a unique doxastic attitude that we are epistemically *obliged* to have given that evidence. We call this thesis the *Uniqueness Thesis* (UT) and its denial the *Permissibility Thesis* (PT). Like many authors writing on this topic, we will interpret the issue *deontologically*: for a belief to be obliged or permitted is for it to be responsible or blameless. Thus, we do not understand the issue merely *evaluatively*, but *deontologically*: the question is not whether doxastic attitudes are obliged or rather merely permitted if they are to be ideally rational or ideally reasonable. The central aim of this paper is to defend PT by rebutting certain arguments against PT and by providing some considerations that favour PT over UT. If we succeed, we have shown that there is good reason to think that responsible or blameless belief is permissible rather than obliged belief. Also, if we succeed, we have shown that the gap between the way we should normatively assess belief and action may not be as wide as has been thought.

This paper is structured as follows. First, we distinguish between several varieties of UT and PT and select those that we deem relevant to our discussion (§ 2). Next, we discuss a couple of objections to PT and argue that they fail (§ 3). However, even if we are correct that they do not succeed, it does not follow that UT is false. Therefore, we also discuss two arguments for PT (§ 4) and provide two motivations for PT by showing that it is entailed by two views that are quite popular among theorists working on doxastic responsibility (§ 5). Even if our strategy does not *establish* or *demonstrate* the truth of PT, it will provide enough reasons to prefer PT over UT.

2. Varieties of PT and UT

Let us start with a word on the terms that we and others working on this issue employ. It seems natural to think that every epistemically obliged belief is an

*This paper is a collaborative effort to which each of us has contributed equally.

epistemically permissible belief. For, it seems paradoxical to say: “I am obliged to stay at home, but I am not permitted to do so,” or “given my evidence, I am obliged to believe that he won the elections, but on the basis of that same evidence I am not permitted to do so.” However, in this paper we use ‘permissible’ as short-hand for ‘*merely* permissible’. Thus, we understand something’s being permitted and something’s being obliged as mutually exclusive states of affairs: if I am obliged to ϕ , then I am not permitted to ϕ , and if I am permitted to ϕ , then I am not obliged to ϕ . We do not think that there is anything unnatural about this use of the word ‘permissible’, for in many conversational contexts, it is used elliptically for ‘*merely* permissible’. If this conversational implicature is clear, then it is *not* paradoxical to say: “She is not *permitted* to stay at home . . . she is *obliged* to do so.” If something is permissible, then it may very well be the only responsible option for *S* (that is, the obliged option), whereas if something is *merely* permissible, then there are at least two options that are equally responsible for *S*. If *S* is neither obliged to stay at home nor obliged not to stay at home, then staying at home and not staying at home are equally permissible for her. Thus, we assume that something is permissible only if it is not obliged and that something is obliged only if it is not permissible. Using the word ‘permissible’ as short-hand for ‘*merely* permissible’ will turn out to be helpful in spelling out the difference between UT and PT.

Before we turn to varieties of UT and PT, we need to make one more preliminary comment. For, what precisely do adherents of UT and PT have in mind when they say that a belief should be obliged or permissible in order for it to be *justified, rational, responsible, reasonable, warranted*, or some such thing? These concepts are notably murky and hard to clarify in a theory-independent way. Given the strongly deontological terms used by several adherents of UT and PT, some of which will be displayed below, it seems right to interpret words like ‘rational’, ‘justified’, and ‘responsible’ deontologically as well. And that means that *S* justifiedly, rationally, or responsibly believes that *p* iff *S* is epistemically blameless¹ for believing that *p*, where one is epistemically blameless for believing that *p* iff one is not the proper object of the negative reactive attitude² of blame for believing that *p*.³ Thus, we will not be discussing any views that cast uniqueness and permissibility in terms of ideally, fully, or perfectly rational (reasonable) doxastic attitudes.⁴

1. The dominant deontological view of justification is that justified belief is *blameless* belief. A handful of philosophers claim that justified belief is *praiseworthy* rather than merely *blameless* belief (e.g. Weatherson 2008). We do not have space enough to deal with this issue here. For an argument against the thesis that justified belief is praiseworthy belief, see Booth and Peels (2010).
2. For the notion of ‘reactive attitude’, see Strawson’s well-known 1974 essay.
3. Below, we discuss some arguments against PT provided or at least inspired by Roger White. It seems plausible to interpret his concept of doxastic justification deontologically. (We do not, however, claim that it *must* be interpreted in this way.) White (2005) uses the words ‘rational’ (pp. 445, 448), ‘reasonable’ (pp. 445, 448), and ‘responsible’ (p. 451) interchangeably, while he uses the word ‘irresponsible’ (p. 448) as the opposite adjective. Some of our arguments may also apply to PT and UT interpreted as theses about perfect rationality, but since perfect rationality is not our concern, we will not flag up when.
4. For two philosophers who understand UT and PT as theses concerning ideal, perfect, or full rationality, see Conee (2010, 70) and Kelly (2010, 120).

Now, what precisely does UT amount to? In order to find out, let us first consider Bruce Russell's statement of it:

A difference between epistemic and moral duty is that one *must* either believe a proposition, disbelieve it, or suspend judgment, while sometimes one can permissibly either perform or not perform some action. There is no epistemic attitude that it is merely permissible to have, while it can be merely permissible to perform some action. If my evidence for and against extra-terrestrial intelligence has equal weight, I must suspend judgment; I cannot suspend judgment, believe, or disbelieve. However, assuming that the moral reasons for and against my now getting up and getting a drink of water are equal, it is permissible for me to do either. It is not the case that I *must* do one or the other.⁵

Russell's view seems to be that for every person *S* and proposition *p*, given the evidence that *S* has concerning *p*, there is a single doxastic attitude that *S* has an epistemic obligation to have towards *p*. This suggests the following understanding of UT:

UT₁: For any person *S*, evidence base *E*, and proposition *p*, there is a unique doxastic attitude (belief, disbelief, or withholding) that, given *E*, *S* is obliged to have towards *p*.

UT₁, however, is not the only version that is available to the adherent of UT. For, many epistemologists are convinced that belief is something that comes in degrees. And this means that an adherent of UT could say that one's evidence determines a unique responsible doxastic attitude but *not* a unique degree of that attitude. Another option would, of course, be to take the strong position that one's evidence does not only determine which doxastic attitude is uniquely responsible, but also which *degree of* belief or disbelief (or withholding if withholdings come in degrees) is uniquely responsible. Adherents of the previous option could stick with UT₁, whereas adherents of the second view could opt for the following version of UT:

UT₂: For any person *S*, evidence base *E*, and proposition *p*, there is a unique doxastic attitude and a unique degree *x* of that attitude⁶ such that, given *E*, *S* is obliged to have that attitude to degree *x* towards *p*.⁷

5. Russell (2001, 36). Cf. Feldman (2007, 205).

6. It is widely believed that, in opposition to belief and disbelief, withholding does not admit of degrees. If this is correct, then UT₂ should be read as follows: For any person *S*, evidence base *E*, and proposition *p*, there is a unique doxastic attitude and, *in the case of belief and disbelief*, a unique degree *x* of that attitude such that, given *E*, *S* is obliged to have that attitude and, *in the case of belief and disbelief*, that attitude to degree *x* towards *p*. Since the current version of UT₂ is much more economical, we will stick to it.

7. One might suggest that there is a third option, to wit the view that one's evidence determines unique responsible degrees of doxastic attitudes, but not which doxastic attitude is responsible. In principle, this is indeed an option, but we are fairly confident that nobody has taken this position. For, it is hard to see how it could be that, say, believing that *p* to degree *x* is responsible, withholding that *p* or disbelieving that *p* to some degree *y* is also responsible, but believing that *p* to a degree *x'* slightly lower than *x* is irresponsible.

There are quite a few epistemologists, though, who say things that clearly conflict with UT. As Matthias Steup rightly remarks,

[a]ccording to a venerable tradition in epistemology, epistemic justification is to be understood as epistemic permissibility. Our beliefs are justified if, and only if, what we believe is epistemically permissible for us to believe.⁸

William Alston formulates the position as follows:

[t]o say that *S* is justified in believing that *p* at time *t* is to say that the relevant rules or principles do not forbid *S*'s believing that *p* at *t*. In believing that *p* at *t*, *S* is not in contravention of any relevant requirements. Again, it is not to say that *S* is required or obligated to believe that *p* at *t*, though that might also be true.⁹

This view entails that UT is false. On the position formulated by Alston, to say that *S* justifiedly or responsibly believes that *p* is to say that *S*'s belief that *p* is *permissible* rather than *obliged*. However, like UT, PT can be spelled out in different ways. First, PT can be taken to apply to *all* or only to *some* cases. According to the former view, for any person *S*, proposition *p*, and evidence base *E*, at least two doxastic attitudes towards *p* are equally permissible for *S*. The latter version of PT confines itself to only *some* persons, only *some* propositions, or only *some* evidence, or displays a combination of these restrictions. Second, PT can range over *all* three doxastic attitudes—belief, disbelief, and withholding—or over only two. In the latter case, the attitudes in question would presumably be those of belief and withholding or those of disbelief and withholding. For, it is hard to see how belief and disbelief could be equally permissible without suspension of judgement being permissible. Third, like UT, PT can cover merely degrees of doxastic attitudes, or both degrees of doxastic attitudes and doxastic attitudes simpliciter.¹⁰

Combining these options would lead to a large number of different versions of PT. Rather than boring the reader by spelling them all out, we present those varieties of PT that play a pivotal role in our discussion below:

PT₁: For any person *S*, evidence base *E*, and proposition *p*, there are at least two doxastic attitudes that, given *E*, *S* is permitted to have towards *p*.

PT₂: For any person *S*, evidence base *E*, and proposition *p*, there are at least two different degrees of some doxastic attitude such that, given *E*, *S* is permitted to have that attitude towards *p* to each of those degrees.

8. Steup (2000, 25).

9. Alston (1989, 116). Cf. Goldman (1986, 70).

10. Again, it is hard to see how one could plausibly defend the third option that one could in principle opt for, namely that different doxastic attitudes are equally permissible, but that different degrees of some doxastic attitude are not, unless, of course, one believes that belief does not admit of degrees.

PT₃: For some persons S , evidence bases E , and propositions p , there are at least two doxastic attitudes that, given E , S is permitted to have towards p .

PT₄: For some persons S , evidence bases E , and propositions p , there are at least two different degrees of some doxastic attitude such that, given E , S is permitted to have that attitude towards p to each of those degrees.

Obviously, PT₁ and PT₂ are stronger than PT₃ and PT₄; they entail, but are not entailed by PT₃ and PT₄. And it seems plausible that if two different doxastic attitudes towards the same proposition are equally permissible, then different degrees of at least one of those attitudes will be equally permissible. If that is correct, then PT₁ is stronger than PT₂ and PT₃ stronger than PT₄. The arguments that we consider in sections 3 and 4 are directed against PT₂ and PT₄. Some considerations in favour of PT that we provide in sections 5 and 6 only count in favour of PT₃, whereas others also count in favour of PT₁.

3. Arguments Against PT

In this section we would like to discuss two arguments against PT that have had considerable influence and that have even been qualified as “formidable.”¹¹ They are put forward by Roger White. We are not sure whether these are really separate arguments. In any case, our response to the first argument will return in our response to the second argument. For each of the two arguments, we argue that it is defective.

1. *The responsiveness to new evidence argument.* The responsiveness to new evidence argument seems directed against PT₄. Since PT₄ is the weakest version of PT that we have given, this argument, if sound, refutes all of PT₁–PT₄. The argument runs as follows. Imagine that two persons S and S^* share total evidence E bearing on p and that S 's degree of belief that p on the basis of E is x , whereas S^* 's degree of belief that p on the basis of E is a lower degree y . Imagine further that S and S^* both acquire the same new evidence E' , so that S^* rationally raises her degree of belief to y' , which happens to be identical to x . Then, there is no good reason for S to raise her degree of belief to a higher degree x' than x , since she is likely to believe that S^* rationally believes that p to degree x on the basis of E & E' and that she already believes that p to degree x . But this, according to White, cannot be correct, for a rational person who believes that p to a degree less than 1, upon acquiring significant new evidence in favour of p , will raise the degree to which she believes that p . If a belief is irresponsive to new evidence, the believer cannot count as rational in holding that belief.¹²

11. See Kelly (2010, 121). White's arguments have recently been criticised by Anthony Brueckner and Alex Bundy (Brueckner & Bundy forthcoming). Unfortunately, we do not have sufficient room to discuss their arguments here. Instead, we will provide some counter-arguments of our own.

12. See White (2005, 454–5).

We are not convinced by the argument. For one thing, it seems that people have different *epistemic standards*, that is, different norms as to what counts as sufficient and relevant evidence for believing, disbelieving, or withholding belief on some proposition p .¹³ If people responsibly hold different epistemic standards, then they might come to have different (degrees of) doxastic attitudes on the basis of the same evidence bearing on p . White responds that one should believe that one's epistemic standards are reliable and that others who come to hold different doxastic attitudes on the basis of the same body of evidence, have epistemic standards that are *not* reliable. According to White, if one does not hold these beliefs, one cannot rationally believe the deliverances of one's own doxastic mechanisms. If PT is correct, White continues, one should also hold that those others rationally adhere to the standards they adhere to. But if that is correct, then it should also seem to one a matter of luck that one has the epistemic standards that one has and, thus, that one is lucky if it turns out that one's epistemic standards are reliable and those of others unreliable. And that would, of course, defeat the rationality of one's belief.

White illustrates his point by giving the following example. Imagine that I know that attending MIT will result in my having certain epistemic standards and attending Berkeley to my having certain different epistemic standards. Imagine also that I believe, as an adherent of PT, that in either case I will hold those standards rationally. Then it should seem to me that it is just lucky if I acquire an epistemic standard that is reliable, and that should defeat my belief's rationality. White even compares such situations to digesting belief-inducing pills. Imagine a situation in which I digest belief-inducing pills. Half of the pills produce true beliefs, whereas the other half produce false beliefs, but I do not know which pills produce true beliefs and which the false ones. Then, White says, the rational attitude towards the propositions involved is that of suspending belief and suspending disbelief. White suggests that an adherent of PT is in a situation relevantly similar to that of someone who has digested belief-inducing pills. For, after having acquired particular epistemic standards, one will believe that one could equally rationally have acquired different epistemic standards that would give different verdicts in the same evidential circumstances.

Is it true that one cannot responsibly continue to adhere to the epistemic standards one has if one believes that others responsibly adhere to different epistemic standards? We think that it is not. Is it a requirement of rationality that one believes that others who do not believe that p have unreliable epistemic standards when they share the same body of evidence bearing on p ? No it is not. One may believe, for instance, that they have reliable epistemic standards but that, apparently, this is one of the relatively few cases in which their epistemic standards do not give the right verdict. As we all know, reliability does not imply infallibility. Thus, one might very well rationally believe that

13. We assume that epistemic standards have propositional content and, hence, are such that there can be evidence for or against them. We take this to be a charitable concession to our opponent, because if epistemic standards were not propositional, then there would be no evidence in favour or against them. And so UT would have the implausible implication that two people with different epistemic standards but the same evidence bearing on p cannot have different but equally responsible doxastic attitudes towards p .

others' epistemic standards are, for all one knows, reliable, but that those others are unlucky in that in this particular case their epistemic standards do not give the right results. In summary, if either (i) S believes or suspends judgement on whether S^* has different, but equally rational epistemic standards for judging the new evidence E' , or (ii) S believes that or suspends judgement on whether S^* has equally reliable epistemic standards that suffer from bad luck on this particular occasion, then S might believe that p to degree x' upon being presented with E' , while also rationally believing that S^* , who also possesses both E and E' , equally rationally believes that p to a lower degree y' .

2. *The argument from self-defeat.* White's second argument against PT is the following. If PT is plausible with regard to many propositions, then it could plausibly be applied to itself. But if it could be plausibly applied to itself, then the evidence that is available to philosophers would not establish whether we should believe PT or rather suspend judgement on it (or even disbelieve it). For, it seems that certain philosophers rationally or justifiedly disbelieve PT. But if one believes that others rationally or justifiedly disbelieve PT on the basis of the same evidence, then it is hard to see how they themselves could (continue to) rationally or justifiedly believe PT.

A minor worry with this argument is that the argument, if convincing, only shows that that one cannot rationally or responsibly adhere to PT. It does not follow that PT is false. Of course, the adherent of PT would in some sense be in trouble if this argument were convincing, for then she could no longer responsibly continue to adhere to and defend PT. But it would still be the case that for all she knows, PT might be true. Thus, this argument would not settle the issue against PT. More importantly, the adherent of PT_3 or PT_4 could very well say that among those propositions to which PT does not apply is PT itself. We do not think that this would be unduly *ad hoc*. It would be so if that were the whole story, if no motivation of this view could be given. But it seems that one *could* motivate this view. One could argue, for instance, that there *is* ample evidence in favour of PT and that everyone who fails to believe that PT is true has not gathered all the relevant evidence, or has not considered the evidence carefully enough, or is irrational in dismissing PT.

But does this argument not at least show that PT_1 and PT_2 are untenable? No, it does not even do that. For, our point about epistemic standards seems equally applicable here. For all we know, opponents of PT may very well have different, but equally responsible epistemic standards, resulting in a different assessment of the same evidence, where that results in a different position on PT. Or, for all we know, opponents of PT may have reliable epistemic standards, as we take ourselves to have, but ones that on this particular occasion suffer from bad luck.

4. Arguments for PT

In this section, we provide two arguments in favour of PT, the first from the impossibility of comprehensive doxastic scrutiny and the second from the indeterminacy of evidential thresholds.

1. *The argument from the impossibility of comprehensive doxastic scrutiny.* Since this argument is quite complicated, let us first re-cap on an issue that came up in the previous section. As we saw, even if two subjects S and S^* share the same evidence E that bears on some proposition p , their epistemic standards for weighing E may differ. Further, since S and S^* may have different evidence in favour of those epistemic standards, they may well come to have different but equally responsible doxastic attitudes towards p . But what if two subjects have the same evidence E bearing on p , but also the same evidence E^* bearing on their epistemic standards? Could S and S^* responsibly take different doxastic attitudes towards p then? If the answer is 'no', then PT_1 and PT_2 are refuted. Because then there would be situations in which a single doxastic attitude is obligated, namely, those situations in which people share the same evidence E bearing on p and the same evidence E^* bearing on the epistemic standards for weighing E . However, unfortunately for the opponent of PT, the answer is 'yes'. Consider how S and S^* might evaluate the evidence E^* in favour of appropriating one epistemic standard and not another. The point is that they may very well *not* have the same epistemic standards for weighing E^* . Thus, they may bring to bear different standards on how to weigh E^* . Here, however, the dangers of either circularity or of engendering an infinite regress present themselves. In other words, S can either enlist his own epistemic standards in assessing the evidence that bears upon the question of whether to take up that standard, but that strategy would be circular. Or, else, S can enlist further beliefs (or standards, if epistemic standards are not beliefs) about the correct way to assess evidence about one's epistemic standards, but to ensure that S and S^* must responsibly come to the same conclusion as regards that evidence, we need to enlist even further beliefs (or standards) about the correct way to assess the evidence that bears on the correct way to assess one's epistemic standards, and so on, ad infinitum.

One reply to the dilemma would be to allow that one's own epistemic standards can legitimately come to bear on the question of how to evaluate E^* in (for example) an effort to secure reflective equilibrium. But this is going to guarantee PT, since S and S^* will always, in theory, be able to appeal to different epistemic standards, such that, in theory, there is always, for any attitude one takes towards a proposition p , a rational alternative one could have taken.

The alternative bullet to bite is to accept that it is not possible for all our beliefs (or epistemic standards) to have been subject to appropriate epistemic checks, but yet that we can be responsible believers even when we believe on the basis of epistemically unscrutinised beliefs. However, this in fact also guarantees PT. Again, suppose that S and S^* are in possession of the same evidence E and E^* . If S and S^* then come to have different doxastic attitudes, should we conclude that one of them is irresponsible? Clearly not, since, if we make the concession outlined above, we can reason that one of S and S^* (or both) based their doxastic attitude on an epistemically unscrutinised belief that happens to be false or a reliable epistemic standard which suffers from bad luck in this particular case.

The opponent of PT may insist that where there is divergence with regard to how S and S^* evaluate E^* , S and S^* do not share the same *total* evidence.

However, as it seems to us, this move is unavailable to the opponent of PT in light of the concessions she must make in order to deal with the problem of facing either circularity or infinite regress mentioned earlier. Both available strategies in effect remove the need to appeal to further evidence E^{**} bearing on the epistemic standards to evaluate E^* . And if we remove that need, then we are in effect saying that the evidence that S and S^* have is identical, *in so far as it is relevant to whether they responsibly believe that p* , even where S and S^* have the same evidence E^* but differ in how they think they should weigh that evidence. That is, there is no evidence E^{**} about which S and S^* differ which is relevant to whether or not they responsibly believe that p . And since there is no way to guarantee that S and S^* ought to concur on what their standards for assessing evidence are, as shown above, then PT, in its strongest mode, must be true.¹⁴

2. *The argument from the indeterminacy of evidential thresholds.* This argument only establishes PT_3 and PT_4 , the weaker versions of the permissibility thesis. It goes as follows. There are cases in which it is indeterminate where exactly the evidential threshold is between whether one ought to have one doxastic attitude towards p and not another. Suppose, for example, that a Spanish teacher, Maria, is testing one of her students, John, on some vocabulary. A year ago she has given John a list of 2,000 Spanish words the meaning of which John is supposed to have now learnt. The exam is an oral one which needs to be completed within a limited amount of time, and she cannot test the student on all the words. How many words does she need to test him on to reasonably believe, and not just suspend judgement on, that the student has learnt all the words on the list? Suppose that she reasonably concludes that once he has got 500 words right, she will consider the student to have successfully completed his task. Imagine, however, that her colleague, Joanna, is also engaged in evaluating the same thing but with another student, James. Imagine that Maria and Joanna know that John and James are equally intelligent and motivated to learn Spanish. Imagine also that Joanna tests James on 490 words. Suppose that both James and John answer all the questions they are asked correctly. The two teachers then both get together to compare their results. Uncontroversially, both will agree that John has passed the test. But is it obvious that they should both suspend judgement on whether or not James has? It does not seem obvious at all. The point here is that in this case, the boundary between having enough evidence to warrant belief as opposed to suspension of judgement is a vague, or indeterminate, one. And if there are cases where that boundary is indeterminate, then there are also cases where two different doxastic attitudes towards p are equally permissible. It seems permissible, for example, for Maria

14. Here, as elsewhere, one may object that this may be true for two different persons, but that we have not been given a reason to think that for *a single person* two different doxastic attitudes towards the same proposition on the same evidence base are equally permissible. We retort that if what we said about two persons is correct, then it also applies to a single cognitive subject. For, a person's blamelessly unscrutinised belief may be true, but it may by bad luck also be false. A person's blamelessly unscrutinised epistemic standard may be reliable and give the correct verdict in a particular case, but one may also suffer from bad luck in that the epistemic standard fails to give the correct verdict in that case. All this is possible while keeping the evidence and the proposition in question constant.

to suspend judgement on whether James has passed the test and permissible for Joanna to believe that he has (or vice-versa) even though both Maria and Joanna share the same evidence, namely, that James knows the meaning of 490 Spanish words on the vocabulary list. Thus there are cases where it is permissible to have either of two different doxastic attitudes towards p on the same evidence base, such that PT_3 (and therefore PT_4) is established. This of course, does not guarantee PT_1 or PT_2 since it does not establish that *all* cases are like this; there may well be cases where the evidential boundary is clear, or where the evidence one already has is well over the threshold, so that, given that body of evidence, only one doxastic attitude is rational. Thus, only our first argument establishes the stronger versions of PT.

5. Motivations for PT

Another way one may be or become convinced of the truth of PT is by noticing that it is entailed by one or several other plausible views that one already adheres to. We will now mention two such views and will show how they entail PT.¹⁵

1. *Responsible peer disagreement.* The first view that we believe entails PT is that there can be such a thing as rational or responsible peer disagreement. That is, that there can be two people who share the same evidence and are considered to be epistemic peers in that they have equivalent intellectual aptitude, experience, etc., and yet both can be considered rational, or reasonable, if they come to diverging conclusions with respect to that evidence. A widely cited example (from Rosen 2001) is that of a jury who has to come to a decision about a difficult case; when the jury members disagree, we seldom take that to mean that someone in the jury is being unreasonable or irrational. The issue of peer disagreement has generated quite some interest in recent years in epistemology. The issue that people seem to have focused on is what the rational outcome ought to be when one is confronted with such a disagreement, that is, on whether we should treat the very fact that there is a disagreement as evidence bearing on what doxastic attitude one ought to take towards a given proposition p . One view, known as the “equal weight view” has it that we ought to revise our attitude in light of the fact that there is a disagreement (i.e. that disagreement counts as evidence), such that both disagreeing parties adjust their doxastic attitudes or confidence in the truth of p and end up no longer disagreeing.¹⁶ The competing view has it that either or both parties need not revise their doxastic attitudes, either not at all or not to the full extent where both parties end up agreeing.¹⁷ The latter, “steadfast” view entails PT and does

15. Roger White mentions several others, but since we are not fully convinced that they do actually entail PT we will not discuss them. In the interest of brevity we will also not explain why we are not. Moreover, since White merely lists these views without fully explaining the entailment, it is difficult to know exactly what we would be arguing against.

16. Cf. Feldman (2006) and Christensen (2007).

17. Cf. Sosa (forthcoming) and Lackey (2010). For a summary and elucidation of the debate, see Christensen (2009).

so even in the steadfast view's weakest mode, that is, as holding that we must make large changes to our doxastic attitudes towards p in light of peer disagreement, but that change need not entail "splitting the difference" between disagreeing parties such that they end up agreeing. If it is rational for two people to come to have diverging doxastic attitudes after they have learnt that an epistemic peer disagrees with them, then it follows that for any attitude one takes towards p there is a rational alternative one could have had given the same evidence, since one can imagine the existence of a hypothetical rational disagreeing peer for any proposition one considers. Thus, it seems that the steadfast view entails the strongest version of PT, PT₁ and, hence, all weaker versions as well.

Now, while the steadfast view entails PT, does the equal weight view entail UT rather than PT? According to Thomas Kelly,¹⁸ it does. However, it seems to us that it does not. First, this is because according to the equal weight view, we only need to revise our doxastic attitudes in cases of what has been called "revealed disagreement,"¹⁹ that is, cases in which we *know* or *believe* that a peer actually disagrees with us. If the view had it that not only believed or known peer disagreement, but unknown peer disagreement as well counted as evidence against p , then the view would implausibly entail that we ought to suspend belief in all those cases in which, unbeknownst to us, some peer would happen to disagree with us. Moreover, it seems that responsibility requires an adherent of this view to suspend judgement on *all* propositions. For, she would believe both that if a peer happens to disagree with her, she ought to suspend judgement, and that, for all she knows, there may be a peer disagreeing with her on *any* proposition she believes. The equal weight view, therefore, has it that only *known* or *believed* peer disagreement calls for suspension of belief. But that means that the equal weight view is silent on UT. An adherent of UT may very well believe that for each proposition, two peers may responsibly take two different doxastic attitudes towards that proposition on the same evidence base, as long as they are not aware of the fact that they disagree with each other. Second, were I not to consider my disagreeing peer's doxastic attitude reasonable (at least before the disagreement became obvious, or if I remained ignorant of the fact that there is disagreement) then I would not rationally consider the fact that we disagree grounds for altering my own doxastic attitude. This means that even the equal weight view gives us strong reasons to believe that at least the weaker versions of PT are true. It implies that we sometimes have good reason to think that there is reasonable peer disagreement, and if there is such a thing, then there is for any doxastic attitude I take towards p a rational alternative, namely that taken by a potential disagreeing peer (though neither the equal weight view, nor the steadfast view have it that one ought to revise one's beliefs in light of that fact). Thus both of the dominant views in the epistemology of disagreement favour PT.

18. More precisely, according to Kelly, an adherent of the equal weight view should endorse UT₂; see Kelly (2010, 119–121).

19. See Feldman and Warfield (2010, 3).

2. *Entailed by doxastic deontologism cashed out in terms of indirect voluntary influence.* We also think that a certain view about what it is to believe responsibly, a view we both countenance, entails PT (recall that the discussion about epistemic obligations and permissions implicitly treats rational belief as responsible belief). The view is the following: in light of the now well-documented claim that belief is not under our direct or long-range voluntary control and the trouble this causes for the possibility of responsible belief in view of the principle that 'ought' implies 'can', we should think that what grounds doxastic responsibility is our ability to influence what we believe.²⁰ To take an analogous example, imagine that I have shy instincts and that when I see an acquaintance pass me on the street I am inclined to look away rather than greet her. Suppose that instinct is so strong in the fleeting moment that I do not have direct control over whether or not I ignore my acquaintance. However, later, and on reflection, I wish I had behaved more politely. And I begin to train myself to automatically reflect on this each time I see an acquaintance and thereby end up behaving in a much less shy fashion. Given that I am able to do this, it seems reasonable to say that, if impoliteness is at least a minimal moral wrong, people who ignore their acquaintances due to shyness can be the proper objects of normative scrutiny, given that there were steps they could have taken to curb their mild misanthropy. Analogously then, since we do not have control over our beliefs, then perhaps the way in which we still can be held to normative doxastic scrutiny is to appeal to the virtuous things we can do to influence our belief-forming habits, such as learning to be more open-minded, learning to be more careful at evaluating evidence, and learning how to spot fallacies of argument and rhetorical tricks. And one may then think that someone is up for normative appraisal on a belief that p if, roughly, it is the case that had she done one of those things she would not have held that belief.

Now, there are some problems with that view.²¹ We think that they can be overcome, but we will not try to defend the view here. For the purposes of this paper, we just want to show that the view entails PT. We think it does for the following reason: suppose that S and S^* responsibly do all the things that influence belief that they should do, such as being careful when evaluating their evidence etc., such that there is nothing either of them could have done that would make them end up with a belief other than the one they have. But suppose that S ends up believing that p , whereas S^* ends up believing that $\sim p$. Perhaps, S 's cognitive abilities differ from those of S^* , perhaps S 's background assumptions are different from those of S^* , perhaps S was just unlucky in that being careful resulted in her having a false belief while S was not unlucky. There are many things that could make it the case that S and S^* end up having different beliefs but both, under the view under discussion, believe responsibly, since there is nothing they could and should have done such that they would have ended up having a different belief. Thus the view entails that it is possible that for any doxastic attitude one takes towards p , there is a rational alternative one could have had (namely, if you are S , that of S^* , and if you are S^* , that of

20. For a defence of this idea, see Nottelmann (2007).

21. For some of those problems, see Alston (1989), Katzoff (2000), and Tidman (1996).

S). It also means that, according to this view of doxastic responsibility, responsible belief is permissible, rather than obliged, belief.

6. Conclusion

In this paper, we have offered two arguments and two motivations for PT. Whereas the argument from the indeterminacy of evidential thresholds and the motivation from responsible peer disagreement counted only in favour of the weaker versions of PT, PT₃ and PT₄, the argument from the impossibility of comprehensive doxastic scrutiny and the motivation from doxastic deontologism cashed out in terms of indirect voluntary influence also counted in favour of PT₁ and PT₂, the stronger versions of PT. We have also shown why the arguments against PT fail. We thus take ourselves to have presented a strong case for why PT should be preferred to UT—at least, in those cases in which PT and UT are understood deontologically rather than merely evaluatively. We, therefore, take it that this tells us something about what it is to believe responsibly, namely that a responsible belief is a permissible and not an obliged belief. In a previous paper, we have argued that responsible belief is blameless rather than praiseworthy belief. The two theses complement each other well, since it seems odd to say that one is worthy of praise each time one does something merely permissible. It seems to us that many deontologists in epistemology have been working with the concepts of permissibility and blamelessness, though we have struggled to find a serious articulation of why they have. These two papers we hope will provide a cumulative case for why they should.²²

References

- Alston, William P. (1989). “The Deontological Conception of Epistemic Justification,” in *Epistemic Justification: Essays in the Theory of Knowledge*. Ithaca/London: Cornell University Press, 115–152.
- Booth, Anthony R. and Rik Peels. (2010). “Why Responsible Belief Is Blameless Belief,” *The Journal of Philosophy* 107.5, 257–265.
- Breckner, Anthony and Alex Bundy. (forthcoming). “On ‘Epistemic Permissiveness’,” *Synthese* 180.
- Christensen, David. (2007). “Epistemology of Disagreement: The Good News,” *Philosophical Review* 116, 187–217.
- Christensen, David. (2009). “Disagreement as Evidence: The Epistemology of Controversy,” *Philosophy Compass* 4, 756–767.
- Conee, Earl. (2010). “Rational Disagreement Defended,” in Richard Feldman and Ted A. Warfield (eds.), *Disagreement*. Oxford: Oxford University Press, 69–90.
- Feldman, Richard. (2006). “Epistemological Puzzles About Disagreement” in Stephen Hetherington (ed.) *Epistemology Futures*. Oxford: Oxford University Press, 216–236.

22. For their helpful comments on earlier versions of this paper, we would like to thank Martijn Blaauw, Catarina Dutilh-Novaes, Jeffrey Glick, Fred Muller, Nikolaj Nottelmann, Herman Philipse, Jeroen de Ridder, Stefan Roski, Conor McHugh, Han van Wietmarschen, and René van Woudenberg.

- Feldman, Richard. (2007). "Reasonable Religious Disagreements," in Louise M. Antony (ed.), *Philosophers without Gods: Meditations on Atheism and the Secular Life*. Oxford: Oxford University Press, 194–214.
- Feldman, Richard and Ted A. Warfield. (2010). "Introduction," in Richard Feldman and Ted A. Warfield (eds.), *Disagreement*. Oxford: Oxford University Press, 1–9.
- Goldman, Alvin I. (1986). *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.
- Katzoff, Charlotte. (2000). "Counter-evidence and the Duty to Critically Reflect," *Analysis* 60.1, 89–96.
- Kelly, Thomas. (2010). "Peer Disagreement and Higher-Order Evidence," in Richard Feldman and Ted A. Warfield (eds.), *Disagreement*. Oxford: Oxford University Press, 111–174.
- Lackey, Jennifer. (2010). "What Should We Do when We Disagree?" in T.S. Gendler and J. Hawthorne (eds.) *Oxford Studies in Epistemology III*. Oxford: Oxford University Press, 274–293.
- Nottelmann, Nikolaj. (2007). *Blameworthy Belief. A Study in Epistemic Deontology*. Dordrecht: Springer.
- Rosen, G. (2001). "Nominalism, Naturalism, Epistemic Relativism," *Philosophical Perspectives* 15, 69–91.
- 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.
- Sosa, Ernest. (forthcoming). "The Epistemology of Disagreement," in Sosa, Ernest. his *Armchair Philosophy*. Princeton: Princeton University Press.
- Steup, Matthias. (2000). "Doxastic Voluntarism and Epistemic Deontology," *Acta Analytica* 15.1, 25–56.
- Strawson, Peter F. (1974). "Freedom and Resentment," in *Freedom and Resentment and Other Essays*. London: Methuen & Co, 1–25.
- Tidman, Paul. (1996). "Critical Reflection: an Alleged Epistemic Duty," *Analysis* 56.4, 268–276.
- Weatherson, Brian. (2008). "Deontology and Descartes' Demon," *The Journal of Philosophy* 105.9, 540–569.
- White, Roger. (2005). "Epistemic Permissiveness," *Philosophical Perspectives* 19, 445–459.