

NORMS OF BELIEF

Mona Simion
KU Leuven and University of St Andrews
Christoph Kelp and Harmen Ghijzen
KU Leuven

Abstract

When in the business of offering an account of the epistemic normativity of belief, one is faced with the following dilemma: strongly externalist norms fail to account for the intuition of justification in radical deception scenarios, while milder norms are incapable to explain what is epistemically wrong with false beliefs. This paper has two main aims; we first look at one way out of the dilemma, defended by Timothy Williamson and Clayton Littlejohn, and argue that it fails. Second, we identify what we take to be the problematic assumption that underlies their account and offer an alternative way out. We put forth a knowledge-first friendly normative framework for belief which grants justification to radically deceived subjects while at the same time acknowledging that their false beliefs are not epistemically good beliefs.

1. Introduction

Justification is widely taken to be normative. The following appears to be an attractive way of capturing this thought:

The Deontic Thesis (DT). One's φ -ing is *prima facie* practically, morally, epistemically, etc. justified if and only if one *prima facie* practically, morally, epistemically, etc. permissibly φ s.¹

If DT captures the way in which justification is normative, then, plausibly enough, the following captures the sense in which epistemic justification of belief is normative:

The Deontic Thesis for Belief (DTB). One's belief that p is epistemically justified if and only if one epistemically permissibly believes that p .²

It may be worth noting that we do not take DTB to offer a substantive account of justification or of the epistemic norm of belief. Rather, we take DTB to capture a highly plausible relation between justified belief and permissible belief. That said, given a substantive account of permissible belief, we can of course use DTB to derive a substantive account of justified belief and vice versa.

Perhaps somewhat surprisingly, DTB quickly leads to a dilemma. To see how, just note that both of the following theses are plausible:

Negative Normative Status of False Beliefs (NFB). False beliefs are epistemically bad and, as such, enjoy epistemically negative normative status. For instance, Faye's belief that Seattle is the capital of France is epistemically bad.

Justification of False Beliefs in Radical Skeptical Scenarios (JFB). The beliefs of agents in radical skeptical scenarios are justified. For instance, when Ned, who has recently been envatted and is now being fed deceptive experiences as of being in a normal physical environment, forms the belief that it's sunny, his belief is justified.

NFB enjoins us to construe the norm of belief as *factive* in the sense that only true beliefs can comply with the norm of belief. After all, if false beliefs have negative normative status, this will have to be because they violate the norm of belief. If false beliefs violate the norm of belief, this means that only true beliefs comply with the norm of belief.

In contrast, JFB enjoins us to adopt an account of justified belief according to which justified belief is *non-factive*. After all, agents like Ned are massively deceived. Nearly all of their beliefs are false. If their beliefs are nonetheless to come out justified, this means that it must be possible to have justified false beliefs.

It is now easy to see how the dilemma arises: If, in view of NFB, we construe the norm of belief as *factive*, it follows from DTB that justified belief must be *factive* also. However, in that case we cannot have JFB. And, conversely, if, in view of JFB, we take justified belief to be *non-factive*, it follows from DTB that the norm of belief is *non-factive* also in which case we cannot hold on to NFB. DTB thus appears to force us to give up at least one of two plausible claims, to wit NFB and JFB. In this way, it leads to a dilemma (which we will henceforth also refer to as 'the permissibility-justification dilemma' or 'PJD' for short).

This paper has two main aims; first, we will evaluate one recent strategy to exit PJD via the first horn, due, respectively, to Timothy Williamson (forthcoming) and Clayton Littlejohn (forthcoming). These authors defend the so-called knowledge norm of belief, and attempt to explain away the intuition of justification for radically deceived subjects in terms of mere blamelessness. We will show that neither of their proposals is ultimately

successful (§§3,4). What's more, §5 will provide independent theoretical reason for thinking that justified belief is non-factive and that radically deceived agents have justified belief. If our argument is successful, it provides reason to favor a non-factive account of the norm of belief. This brings us to the second and main aim of this paper: putting forth a general normative framework that enables us to escape PJD, that is, to allow for justification to be present in demon victims while, at the same time, explaining what is bad about false beliefs (§6). Crucially, the framework does not require us to abandon knowledge first epistemology. On the contrary, it is perfectly compatible with the independently plausible idea that knowledge is the norm of belief in a sense to be specified.

First things first, however, we will now (§2) briefly present the accounts of the norm of belief and of justified belief defended by knowledge firsters such as Williamson and Littlejohn.

2. The Knowledge Norm of Belief

There is a growing number of epistemologists who accept the following norm of belief:

The Knowledge Norm of Belief (KNB). One must: believe that p iff one knows that p .

A belief is permissible by KNB, then, iff it qualifies as knowledge. It is easy to see that, on KNB, the norm of belief is factive. After all, knowledge is factive. If permissible belief requires knowledge, then only true beliefs are permissible. In this way, KNB can easily make sense of NFB.

But now note that, in conjunction with DTB, KNB entails:

Justified Belief is Knowledge (JB = K). One's belief that p is justified if and only if one knows that p

It may be worth noting that a number of knowledge firsters, including Williamson, accept $JB = K$. At the same time, it is easy to see that anyone committed to $JB = K$ will be unable to accommodate JFB. After all, as we already noted, it follows from $JB = K$ that only true beliefs are justified. Since the beliefs of agents in radical skeptical scenarios are nearly all false, they are not justified.

Knowledge firsters typically take a knowledge norm of belief to enjoy independent theoretical motivation. The fact that NFB needs to be accommodated is one of the *prima sources* in this regard. And since, together with the uncontentious DTB, KNB entails $JB = K$, the same goes for $JB = K$. For that reason, knowledge firsters typically think that the right way out

of the PJD is to explain away the intuition of justification in radical deception scenarios. In particular, one popular line has been to argue that, whilst falling short of justification, radically deceived subjects nonetheless enjoy an excellent excuse for believing as they do, as a result of which their beliefs are blameless. According to advocates of this line, those who have the intuition of justification here simply confuse justification with blamelessness. In what follows, we will take a closer look at two prominent defenses of this strategy in recent literature, one due to Williamson (forthcoming), the other due to Littlejohn (forthcoming).

3. Primary and Derivative Norms

3.1. Williamson's normative framework

Williamson's account of the intuition of justified belief in cases of radical deception makes use of a normative framework that is developed in analogy to some clear cases of norm-compliance and norm-violation outside of the realm of epistemology. It will be helpful to start with one of these cases before going on to the application to epistemology.

Following Williamson, let's consider promises. Williamson holds, quite plausibly, that one complies with the norm of promise keeping if and only if one does what one promised to do (forthcoming, p.4). For instance, if you promise to pick up John at the airport by 2pm, then you will have kept your promise if and only if you pick up John at the airport by 2pm. If you are somehow prevented from reaching the airport, say, by car trouble, then you will not have kept your promise. Indeed, you will not have kept your promise even if you have very good reasons for not going to the airport, or if you have done everything in your power to reach the airport but nevertheless failed to reach it. On the other hand, if you try your best to be late but nevertheless arrive at the airport on time, then you will have kept your promise. So one complies with the norm of promise keeping if and only if one does what one promised to do.

Next, Williamson introduces a distinction between what he calls 'primary norms' and 'derivative norms'. One example of a primary norm is the norm of promise keeping just mentioned. Crucially, according to Williamson, justification is linked exclusively to compliance with primary norms. This means that the relevant form of permissibility at issue in DT is unpacked in terms of compliance with the primary norm.

Derivative norms are generated by primary norms. By way of illustration, let's return to the case of promise keeping. The thought is that the primary norm of promise keeping (P) gives rise to a secondary norm (DP) that requires people to have the disposition to keep their promises and a tertiary norm (ODP) requiring people to do what a promise keeper would do in a

given situation. The reason for this is that given that it's good to keep your promises, there also seems to be something good about being the kind of person who is generally disposed to keep their promises, and, furthermore, to do what this kind of person would do. One reason for this is that complying with these derivative norms will usually make it more likely that one also complies with the primary norm.

Of course, since justification is connected to primary norms, compliance with secondary and tertiary norms cannot give one justification. However, Williamson grants that even when one breaks the primary norm, one may still comply with the secondary and/or tertiary norms generated by it. Moreover, and most importantly for present purposes, compliance with secondary and tertiary norms may still make for a good excuse, thus rendering the agent blameless for violating the primary norm. For instance, if you promised to pick up John at the airport by 2pm and arrive there only at 3pm you did not keep your promise. You violated the primary norm of promise-keeping, P. Even so, suppose that the reason why you are late is that you had car trouble, then borrowed a car from a friend to still be able to keep your promise, and then had car trouble again, thereby becoming unable to fulfill your promise. What you did is certainly something a promise keeper would do. As a result you satisfy ODP. If, in addition, you are a promise keeper, you also satisfy DP. In this case, even though you violate P, you are blameless for doing so in virtue of satisfying DP and ODP. In this way, complying with derivative norms may generate excuses for violations of the primary norms.

It's not too hard to see how Williamson wants to apply this normative framework in epistemology, and specifically to radically deceived agents. An important motivation for the idea that radically deceived agents have justified beliefs is that such subjects are "following exactly the same cognitive instincts as [the ones they follow] in forming the same belief in the good case" (forthcoming, p.18). Radically deceived agents thus believe what a good believer—i.e., someone who believes only what they know—would believe in the current situation, and they so believe because they are the kind of people who believe only what they know. To put this in normative terms again, if the primary norm (J) for belief is knowledge, then there is also a derivative secondary norm that one should be generally disposed to believe only what one knows (DJ), and a tertiary norm that one should believe what one who is generally disposed to believe only what one knows would believe in the current situation (ODJ). On this picture, our radically deceived agent, Ned, does not comply with the primary norm for belief. As a result, he is not justified. However, he does comply with the derivative norms, DJ and ODJ. Ned believes exactly what someone who is disposed to believe only what they know would believe in the current situation, thus satisfying ODJ. Moreover, we may assume, he is disposed to believe only what he knows and so complies with DJ. Just as in the aforementioned

case of promise keeping, even though Ned violates the primary norm, J, he has an excellent excuse for this in virtue of the fact that he satisfies DJ and ODJ.

With Williamson's framework in play, let's return to the question as to how he explains the intuition that the beliefs of agents like Ned are justified. The answer is that Williamson offers an error theory for this intuition that has non-epistemic precedent and theoretical backing by the normative framework. To see how this works, let's return once more to the case of promise keeping. In our toy case, the fact that you do what a promise-keeper would do (thus complying with ODP) and, in fact, are a promise keeper yourself (thus complying with DP) renders you blameless here and may lead us to sympathize with you. When presented with the question as to whether you are justified in acting as you did, our sympathies may incline us to think that you were justified rather than unjustified. However, this does not change the fact that you violated P, the primary norm of promise keeping. Since justification is tied to compliance with the primary norm, this means that your action is really unjustified. Rather, what is going on with those who have the intuition of justification is that they mistake blamelessness for justification. And the same goes, *mutatis mutandis*, for the beliefs of radically deceived agents. The fact that Ned is someone who is generally disposed to believe only what he knows and that he believes what a person so disposed would believe (thus complying with DJ and ODJ) renders him blameless and may lead us to sympathize with him. When presented with the question as to whether Ned's beliefs are justified, it may even lead us to think that they are justified rather than unjustified. But just as in the case of promise keeping, this simply does not change the fact that Ned violated J, the primary norm of belief, and that, as a result, Ned's beliefs are really unjustified. Again, those who have the intuition of justification mistake blamelessness for justification.³

Williamson thus exits the dilemma by embracing a factive norm of belief and offering an error theory for the intuition of justification in cases of radically deceived agents.

3.2. Against mere blamelessness

While Williamson's treatment of the dilemma is attractive, there is also a problematic feature of his framework. There is reason to believe that Williamson's error theory remains ultimately less than fully satisfactory. To see this, consider the following case:

Brainwashed Cognizer. Unfortunate Brianna has recently been brainwashed into taking a red sky in the evening as an indication that something bad will happen to her soon. Just now Brianna has noticed that the evening sky is red and has formed the corresponding belief about her near future.

Intuitively, Brianna's belief that something bad will happen to her soon is not justified. It is easy to see that Brianna does not comply with a truth-related norm J (suppose, as we may, that nothing bad will happen to Brianna). Even so, she does comply with derivative norms DJ and ODJ, or so we may assume. After all, Brianna may well be a person who is generally disposed to believe only what she knows (including about her future), and she believes what someone who is so disposed (i.e., she herself) would believe in the current situation.⁴ But now notice that this means that Brianna's belief is on a normative par, as it were, with the beliefs of Ned, at least as far as Williamson's framework is concerned. After all, both violate the primary norm of belief, J, whilst complying with the derivative norms, DJ and ODJ. Moreover, it is independently plausible that, in both cases, the agents are blameless for believing as they do. Crucially, however, while the case of Ned generates an intuition of justification, in the case of Brianna we get just the opposite intuition. So what Williamson's framework has not yet explained is why only some of the agents who comply with DJ and ODJ appear justified. A factor that is at least relevant to our intuitions about justification is still left out of the picture.⁵ By the same token, the worry remains that this difference in intuitions is sourced in another relevant normative distinction between Brianna and Ned, one that is not captured by Williamson's framework.

One response to this complaint is that there are simply different ways of being blameless and different ways of being excused. Given the variety of excuses that people use in response to norm violations, one should expect there to be different kinds of blamelessness. Our intuitions might then simply be tracking these different kinds of blamelessness rather than a difference between being blameless and being justified. But to make this response stick one should be able to offer an account of at least some of these different kinds of blamelessness. In the following section, we'll turn to Littlejohn's proposal, which does exactly this.

4. Justifications, Excuses and Exemptions

3.1. *Littlejohn and the trichotomous scheme*

According to Littlejohn (forthcoming), we should hesitate to draw conclusions about a difference in justification on the basis of an intuitive difference between the case of Ned, our radically deceived agent, and the case of a brainwashed subject like Brianna:

Why shouldn't we say that this [Ned's case] is just another case of excuse, albeit one in which the grounds for excusing differ from those that are present in cases of brainwashing? If the category of excuse is sufficiently heterogenous, we could say that insanity and ignorance both function as excuses, albeit excuses that do

their work in different ways. Maybe it's right to say that mere blamelessness doesn't account for the positive reaction to the case, but maybe it's a mistake to equate excuses with mere blamelessness. All we know about excuses thus far is that they aren't justifications and if the category isn't homogenous, maybe some excuses involve some sort of positive element, something that we might admire, praise, commend, etc. (forthcoming, p.9).

Inspired by Strawson (1962), Littlejohn ventures to work out this idea in more detail by invoking a trichotomous normative framework, in which excuses and justifications are distinguished from a third category, i.e., exemptions (forthcoming, p.10). The thought is that subjects are exempt from norm compliance when their rational capacities are absent or compromised to such an extent that they simply do not have the required prerequisites for being responsible agents who can be held accountable for their actions. This kind of blamelessness is different from one where an agent has an excuse for their actions: in such a case all the required rational capacities are present, they are exercised in an excellent way, yet a norm is still violated. Indeed, if one didn't exercise one's rational capacities in an excellent way, then one wouldn't be excused for violating the norm. Just suppose that the reason why you didn't pick up John at the airport by 2pm is that you simply forgot about it. Forgetting something that you should have taken care to remember is exactly not the sort of excuse that will get you off the hook for breaking one's promise. What's needed to excuse oneself for norm violation is an excellent exercise of one's rational capacities.

Both these forms of blameless norm-violation, i.e., exemption and excuse, should be distinguished from genuine compliance with the relevant norm and hence justification. (Note that Littlejohn, like Williamson, ties justification to compliance with what Williamson would call the relevant primary norm.) It is only if one actually fulfills the conditions of one's promise, and does what one promised to do, that one has complied with the norm of promise keeping and so acts justifiably. Similarly, it is only when one knows what one believes that one has complied with the knowledge norm of belief, and so believes justifiably.

The threefold distinction between exemption, excuse, and justification appears to provide us with what we need to make sense of the intuitive difference between Brianna and Ned. Whereas Brianna's rational capacities are momentarily compromised by having been brainwashed into taking a red sky in the evening as an indication that something bad will happen to her soon, Ned has no such difficulties. Ned's rational capacities are intact, and, in fact, one could argue that Ned exercises those capacities in an excellent way in coming to believe as he does. Ned is forming beliefs on the basis of his perceptual experiences in a way that would have provided him with accurate beliefs about his environment if only he had not been radically deceived. This means that Ned is excused for believing as he does, whereas Brianna

is exempted. The difference between excuse and exemption thus allows us to explain the intuitive difference between the cases of Brianna and Ned without supposing that either of them has justified beliefs. Yes, radically deceived subjects are exercising their rational capacities in an excellent way, but this could just show that they have an excellent excuse rather than a fully-fledged justification. More is needed to show otherwise.

3.2. *Against mere excuses*

Unfortunately, there is reason to believe that Littlejohn's account also remains ultimately unsatisfactory. To see why, consider first the following case (adapted from Goldman 1988):

Benighted Cognizer. Ben is a member of an isolated and benighted community.

Many of his methods of belief formation have no connection to truth whatsoever, but they are common lore in Ben's community. Let's suppose that Ben wants to know the best time to sow his crops. According to the lore of his community, in order to achieve this, he will first have to sacrifice a goat and bury it in a sacred place. Then he must sit outside his house until it starts to rain and then return to the burial place. If the sun is shining again by the time he will have arrived, it is time to sow the crops. If not, he will have to return home and continue sitting outside his house until the next rainfall. Ben has flawlessly implemented this procedure and has thereby arrived at a belief that it is time to sow the crops.

Ben's belief is unjustified. After all, Ben's belief is the outcome of a method that is, unbeknownst to him, an epistemically disastrous way of forming beliefs. Even so, Ben is surely still blameless for not knowing this, given that the method is so widely accepted within his own community.

Crucially, in this case, Ben's rational capacities do not appear compromised in a way that makes him unaccountable for his actions; instead, Ben appears to exercise his rational capacities in an excellent way given his circumstances. Ben does what most rational people would do in his situation: to find out the answer to his question, he uses a complex method that enjoys widespread acceptance within his community. Ben's case, like Ned's, thus falls in Littlejohn's category of excuse: Ben is blameless for believing as he does because, although he has violated the primary norm of belief, he has exercised his rational capacities excellently.

This means that there is, again, an intuitive difference between two cases—i.e., the case of Ned and the case of Ben—that is not explained by Littlejohn's trichotomous framework. Ned but not Ben appears to be genuinely justified. However, by Littlejohn's framework, both Ben and Ned end up on a normative par: both fail to satisfy the norm of belief and both are blameless in virtue of enjoying an excuse for their beliefs. Again, there

seems to be a relevant distinction, this time between Ben and Ned, which is not captured by the proposed framework.

Of course, one could now try to argue that there is yet another distinction to be made between types of excuses, where Ben falls on one side of the distinction, and Ned on the other. However, we think that the prospects for such a move are dim. The reason for this is that alongside these intuitive difficulties, there is also theoretical reason to think that the Williamson-Littlejohn line is mistaken. To see this, note first that there is an important structural difference between the cases of Ned on the one hand and of Ben and Brianna on the other. Ned forms his beliefs in ways that enjoy a strong connection to truth and knowledge: in normal circumstances, his ways of belief formation do give him true beliefs and even knowledge. In fact, there is a clear sense in which he is simply unlucky to be in the predicament he finds himself in and so unlucky not to acquire true beliefs and knowledge. In contrast, the same is not true of either Ben or Brianna. For them it is not as if had circumstances be more hospitable, their ways of belief formation would have generated true belief, never mind knowledge. Rather, Ben and Brianna's ways of belief formation enjoy no connection to truth and knowledge at all.

Second, there is reason to believe that this structural difference is a difference worth marking in epistemology. In fact, it is a difference worth marking as a distinctively normative difference. To see this, note that as epistemic agents we frequently ask the same or at least similar questions. Moreover, we use the same ways of belief formation to answer these questions. Given that this is so, however, it makes sense to accord positive normative status to beliefs that are formed by ways of belief formation that enjoy a strong tie to truth and knowledge. After all, so doing will reinforce the epistemic behavior of those using these ways of belief formation. By granting Ned positive normative status for his beliefs, we are reinforcing his use of his ways of belief formation. Moreover, we also encourage others who might be using the same ways of belief formation to continue doing so. Given their strong tie to truth and knowledge, this clearly is a good idea.

Correlatively, it also makes sense to accord negative normative status to beliefs that are formed by ways of belief formation that do not enjoy any connection to truth and knowledge as so doing will incentivize those using such ways of belief formation to stop doing so. By according Brianna and Ben's beliefs negative normative status, we are criticising them for forming beliefs in the way they do and we signal to others not to adopt such ways of belief formation. Given that these ways of belief formation enjoy no connection to truth and knowledge, this is also clearly a good idea.

Finally, while it makes sense to accord negative normative status here, it also makes sense to grant Brianna and Ben an excuse. After all, their use of that way of belief formation was out of their control (Brianna) or else done in good conscience (Ben).

It thus comes to light not only that there is an important structural difference between the cases of Ned on the one hand and Ben and Brianna on the other, but also that it makes sense to mark this difference as a distinctively normative difference. In particular, while it makes sense to grant Ned's beliefs positive normative status, the same is not true of Ben's and Brianna's beliefs, to which it makes sense to accord negative such status, albeit with an excuse. It should now be clear that those pursuing the Williamson-Littlejohn line will be hard pressed to accommodate this distinction. Since it is a distinction it makes sense to accommodate, this is bad news for them.

4. A New Normative Framework

In the beginning of this paper, we identified a dilemma according to which we either adopt a factive norm of belief in which case we can explain what's bad about false beliefs but won't be able to allow radically deceived agents to have justified beliefs, or else we adopt a non-factive norm of belief in which case we may just be able to accommodate radically deceived agents but will no longer be able to explain what's bad about false beliefs.

In the meantime, we have seen that there is intuitive as well as theoretical reason to accord the beliefs of radically deceived agents positive normative status. By the same token, there is reason to believe that they satisfy the norm of belief. It looks as though the proper way out of the PJD is by opting for a non-factive norm of belief and loosening up the tie between justification and truth. Of course, this means that we are still facing the question of why false beliefs are bad. This section attempts to offer an independently motivated answer to this question that simultaneously allows us to hold on to DTB and to grant justification to radically deceived agents.⁶ In a nutshell, what we will try to argue is that the above is actually a false dilemma, resting on a mistaken normative assumption.

4.1. *Prescriptive and evaluative norms*

First, let us have a look at a useful distinction made by Conor McHugh (2012). According to McHugh, one can distinguish between evaluative and prescriptive norms. Prescriptive norms are primarily about what one ought to do; they prescribe certain pieces of conduct for agents (McHugh 2012, p.9).⁷ Paradigmatic examples of prescriptive norms include moral norms such as 'Don't steal' or 'Don't lie' but also traffic norms such as 'Drive 50km/h within city bounds' and rules of games such as 'only move the bishop diagonally' in chess. We take it that one crucial function of prescriptive norms is to reinforce certain forms of conduct (permitted ones) and to discourage others (prohibited ones). For instance, one function of the traffic norm that

requires people to drive no faster than 50km/h within city bounds is to reinforce staying within the 50km/h speed limit and to discourage speeding.

In contrast, evaluative norms regulate what it takes for a token of a particular type to be good or bad with regard to its type. Take, for instance, the norm that a good hospital is a clean hospital, that a good knife is sharp or that good driving is safe driving. Evaluative norms use 'good' in Geach's (1956) attributive sense (McHugh 2012, p.22), where 'good' functions as a predicate modifier, rather than as a predicate in its own right. When the evaluative norm states that good knives are sharp, it merely states that knives *qua knives* are good only if they are sharp. It does not entail that good knives are good simpliciter,⁸ or good for some purpose or another. Similarly, it might be true that good burglars are stealthy (stealth thus being a condition on complying with the evaluative norm for burglars), but this does not entail that good burglars are good simpliciter.

Evaluative norms differ from prescriptive norms in that they don't prescribe a certain piece of conduct for agents. This is not to say that evaluative norms do not feature genuine oughts. On the contrary, there is a clear sense in which evaluative norms tell us something about how things ought to be. For instance, the aforementioned evaluative norm for hospitals tells us that hospitals ought to be clean, the norm for knives that knives ought to be sharp and the norm for driving that driving ought to be safe. But this does not detract from the fact that such norms still do not prescribe a certain course of conduct and, as a result, are not prescriptive norms. The main difference lies in the fact that, while prescriptive norms are ought-to-dos, evaluative norms are ought-to-bes.

Evaluative and prescriptive norms can come apart. It is entirely possible for an evaluative norm to be violated without a prescriptive norm being violated. Consider for instance the following evaluative norm: good sunsets are not too cloudy. On a very cloudy evening this evaluative norm will not be satisfied. Compatibly with that, no prescriptive norm may be violated. And, conversely, it is possible to violate a prescriptive norm without violating any evaluative norm. Suppose your boss has imposed a completely pointless rule for making coffee according to which coffee powder must be scooped into the filter alternately with a blue and a red teaspoon. Suppose you violate this rule, say because you are only using the red spoon. You have broken a prescriptive rule for coffee making at your office. However, it seems plausible that you need not have violated some evaluative norm as well. In particular, the coffee that you are making may still be good coffee, the way in which you are making it may still be a good way of making coffee and you may still be a good barista.

While evaluative norms thus differ from prescriptive norms, the two may still be related. In particular, prescriptive norms often enough derive from evaluative norms. They serve to ensure that the evaluative norm is likely enough complied with. For instance, prescriptive norms of driving, such as the norm 'Drive no more than 50km/h within city bounds', serve to ensure

that the evaluative norm of driving—according to which good driving is safe driving—is likely enough complied with. In this way, evaluative norms often come first and prescriptive norms are in their service.

4.2. *Norm coincidence*

With the distinction between prescriptive and evaluative norms in play, let's return to the case of belief. In particular, let's ask whether factive accounts of the norm of belief should be considered evaluative or prescriptive norms. For simplicity, let us focus on the knowledge norm: is the norm an evaluative one, i.e. stating that a good belief qualifies as knowledge? Or is it a norm prescribing how one is to go about belief formation, i.e. prescribing that one should only hold knowledgeable beliefs? One answer to this question is suggested in the following passage from Williamson:

“If justification is the fundamental epistemic norm of belief, and a belief ought to constitute knowledge, then justification should be understood in terms of knowledge too. Indeed, a belief is fully justified if and only if it constitutes knowledge” (2014, p.5).

It would seem that what Williamson has in mind is an evaluative norm, stating that what it takes for a belief to be a good belief is for it to be knowledge. After all, what Williamson says here is that a belief ought to be knowledge and we saw that ought-to-bes are the stuff that evaluative norms are made of.⁹ It appears, then, that Williamson is attracted to the following:

The Evaluative Knowledge Norm of Belief (EKNB). One's belief that *p* is a good belief iff one knows that *p*.

But now recall that Williamson also accepts:

The Knowledge Norm of Belief (KNB). One must: believe that *p* iff one knows that *p*.

Note that KNB is an ought-to-do rather than an ought-to-be; that is, a norm that prescribes a certain piece of conduct for agents, and is an apt guide for conduct. KNB tells us how one must go about believing. Unlike EKNB, KNB is therefore a prescriptive norm. It is easy to see that EKNB and KNB jointly entail:

The Coincidence Thesis for Belief (CTB). A belief is good (i.e. satisfies the evaluative norm of belief) iff it satisfies the prescriptive norm of belief.

Now, we do not mean to suggest that there is a principled reason why in the case of belief the evaluative and prescriptive norms couldn't coincide. However, as the discussion in the previous subsection indicated, there also

is little reason to think that the two must go hand in hand. Furthermore, one serious downside of CTB is that it leads straight to the permissibility-justification dilemma. After all, if the evaluative and the prescriptive norms for belief coincide, they will either both be non-factive, and thus fail to explain what is wrong with false beliefs, or they will both be factive and thus, in conjunction with DTB, will deny justification to radically deceived subjects. Another downside, which specifically affects champions of a knowledge norm of belief, is that it forces them to exit PJD via the first horn, i.e. by denying justification to radically deceived agents (at least given DTB). And we have already seen that there is reason not to do so.

We believe that PJD is ultimately rooted in CTB. Once we abandon CTB and keep the evaluative norm of belief neatly separated from its prescriptive cousin, a way out of PJD opens up. This way is attractive in that it not only allows us to accommodate both NFB and JFB, but is also backed by an independently plausible normative framework and is available to knowledge firsters. In the remainder of the section, we will sketch how it goes.

4.3. Our alternative

We have seen that, generally speaking, evaluative and prescriptive norms can come apart in the sense that it is at least in principle possible to violate a certain evaluative norm without violating any prescriptive norm and vice versa. The key idea of our proposal is to allow that, in the case of belief, the evaluative and prescriptive norms can come apart also. In particular, it is possible to violate the evaluative norm of belief whilst complying with the prescriptive norm. This will allow us to adopt an evaluative norm of belief that is factive whilst leaving room for a prescriptive norm that isn't. It is easy to see that this opens up the possibility of exiting PJD in a way that accommodates both NFB, that there is something bad about false beliefs, and JFB, that radically deceived agents are still justified. After all, given the factive evaluative norm, false beliefs will violate this norm and so won't be good beliefs. And, if we take 'permissibility' in DTB to be permissibility by the non-factive, prescriptive norm, there is no longer any in-principle obstacle to allowing that radically deceived agents believe justifiably.

To see why this makes good normative sense, recall first that section 3.2 offered a detailed argument in favor of according positive normative status to the beliefs of agents in radical skeptical scenarios. This was because their ways of belief formation are connected to true belief and knowledge and that it makes sense to reinforce use of such ways of belief formation. But now recall also that a key function of prescriptive norms is precisely to reinforce certain forms of conduct by permitting them. The normative profile of the beliefs of radically deceived agents (as per section 3.2) is thus just what we would expect it to be were these beliefs in compliance with the prescriptive

norm of belief. One highly attractive explanation of this fact is that these beliefs do indeed satisfy the prescriptive norm of belief. We thus have good abductive reason to think that radically deceived agents do indeed comply with the prescriptive norm of belief.

At the same time, second, it remains highly plausible that there is something bad about false beliefs, i.e. that they have negative normative status. If so, it also remains plausible that there is some norm of belief that is violated by false beliefs. Of course, if even radically deceived agents satisfy the prescriptive norm of belief, the norm of belief that false beliefs violate cannot be the prescriptive norm. So, it has to be a different kind of norm. Our proposal that it is the evaluative norm not only answers the question as to which kind of norm this might be, it also fits the original intuition that there is something *bad* about false beliefs.

Moreover, our proposal is open to knowledge firsters. To see this, note first that knowledge firsters are free to adopt an evaluative norm of belief according to which good belief is belief that qualifies as knowledge. This leaves the question as to whether the prescriptive norm can also be understood in a knowledge first friendly fashion. Again, the answer is yes. To see this, suppose that, even though the evaluative and prescriptive norms of belief come apart, they are also related in much the same way in which we found these kinds of norm to be related in the traffic case: the prescriptive norms serve to ensure that the evaluative norm is often likely enough satisfied, at least in suitable circumstances. If so, the prescriptive norm of belief is derivative from the evaluative norm. And, since the evaluative norm of belief is knowledge, there is still a clear sense in which knowledge comes first in this framework. What we will have, then, is an account of the normativity of belief that avoids PJD, is available to knowledge firsters, and enjoys backing by an attractive normative framework.¹⁰

Now, one might worry whether there actually is some substantive account of the prescriptive norm of belief that is both knowledge first and fits the above description in the sense that the prescriptive norm is non-factive and can be viewed as derivative from the evaluative norm in the way envisaged. Fortunately, there are a number of accounts that fit the bill. For present purposes we will focus on two here, each one has been advocated by some of us elsewhere.

The first is Simion's (2016) knowledge first functionalism. The account is proper functionalist in that it takes the epistemic normativity of belief to drop out of the epistemic function¹¹ of our cognitive systems. It is knowledge-first epistemological in that, unlike traditional proper functionalism, it unpacks the function at issue here in terms of knowledge. Roughly, the thought goes as follows: just as the economic function of knife producing systems is to produce good, sharp knives, the epistemic function of our belief forming systems is to produce good, knowledgeable beliefs. Mere belief, then, is a failure on the part of our cognitive system to fulfill its epistemic function just as blunt

knives are failures on the part of knife producers to fulfill their economic function. In turn, functions give rise to prescriptive norms: a token of type T with function F will be functioning properly when functioning normally, i.e. in the way in which it did back when it acquired its function. Knife producing systems will be functioning properly when functioning normally, that is, in the way in which they did when they acquired their economic function of producing sharp knives. Similarly, our cognitive systems will be functioning properly when functioning normally, that is, in the way they functioned back in the day when they acquired the function of generating knowledge.

On this view, then, epistemic justification consists in the proper performance of cognitive systems that have generating knowledge as their epistemic function. The standards for normal functioning are thus constitutively associated with promoting knowledgeable beliefs.

Of course, normal functioning need not imply function fulfillment; knife producers can function normally—i.e., follow the normal production steps—and still fail to fulfill their economic function of producing sharp knives, due to, say, unfriendly environmental conditions. Similarly, our cognitive systems can function normally, but still fail to produce knowledgeable beliefs. It is easy to see that, on this account, Ned's beliefs come out as justified, for his cognitive systems, the ones in charge with generating knowledgeable beliefs, are functioning normally and, as such, comply with the epistemic prescriptive norm (cf. Graham 2012 for a similar account within the traditionalist framework). At the same time, however, Ned's cognitive system fails to fulfill its epistemic function of producing knowledgeable belief. As a result, Ned's beliefs are not good beliefs.

Let's then move on to the second view, which is the knowledge first version of virtue epistemology, defended in Kelp (forthcoming-*a,b*) and Kelp and Ghijssen (2016). According to this view, inquiry is taken to be a goal-directed practice with moves and success, like target archery. In target archery, moves are shots and successes hits. In inquiry, beliefs are moves and knowledge is success. Besides success, moves in such practices can be assessed as competent and apt. In particular, a move is competent iff, roughly, it is the result of an ability to produce successful moves. In the case of inquiry, this means that a competent belief is, roughly, belief that is produced by an ability to know. One key thesis is that a belief is justified iff it is competent in this sense. By DTB, this means that a belief satisfies the prescriptive norm of belief iff it is competent.

Now, moves in general can be competent without being successful. One can fire a competent shot that misses the target, thanks to a gust of wind, say. Belief makes no exception here. This means that it is possible to have justified beliefs that fall short of knowledge. What's more, as Kelp (forthcoming-*b*) and Kelp and Ghijssen (2016) argues in some detail, this view also allows for justified false belief. In fact, the view grants agents in radical skeptical scenarios justified belief (Kelp and Ghijssen 2016, p.181, n.14). Given DTB,

we get a substantive account of the prescriptive norm of belief according to which this norm is not factive and is satisfied by radically deceived agents. JFB is accommodated.

Finally, let's turn to the question of the relation between the prescriptive and the evaluative norm. To answer it, let's take another look at the phenomenon of a competent shot in target archery. Why do we accord positive normative status to such shots? Very plausibly the answer is that competent shots are produced by an ability to hit the target and that such an ability makes successes highly likely.¹² Given that this is so, there is reason to believe that the positive normative status of a competent shot in target archery derives from the positive normative status of a success in that practice. After all, competent shots enjoy positive normative status because they are produced in a way that makes successes likely, at least in favorable circumstances. *Mutatis mutandis*, in the case of belief we get the result that competent beliefs enjoy positive normative status because they are likely to be successful—i.e. to qualify as knowledge—at least in favorable circumstances. Since, on the present view, one complies with the prescriptive norm of belief just if one believes competently and with its evaluative cousin just if one knows, this means that the prescriptive norm of belief is derivative of the evaluative in just the way envisaged.

5. Conclusion

We have argued that the beliefs of radically deceived subjects appear to have an epistemic status that is better than mere blamelessness or mere excuse. What's more, the idea that false beliefs can still be justified makes sense both from the way in which norms work in general, and from a social perspective in which one wants to reinforce ways of belief-formation that are strongly connected to truth and knowledge. These considerations cannot be accommodated by either Williamson's or Littlejohn's normative framework, in which knowledge remains the norm of belief, and false beliefs can at best be excusably held. In contrast to these theories, we have proposed a framework in which knowledge is taken as the evaluative norm of belief, and justification arises from compliance with an associated prescriptive norm of belief where prescriptive norms have the function of ensuring that the corresponding evaluative norms are often enough satisfied. One can thus uphold a genuinely knowledge-first framework of knowledge and justification without upholding the idea that false beliefs cannot be justified.

Notes

1. For the purposes of this paper, we will bracket the phenomenon of defeat. For that reason we will henceforth take the *prima facie*-proviso as read.

2. When it comes to belief, we will only discuss epistemic justification. As a result, we will use 'justification' as a shorthand for 'epistemic justification' in the context of attributions of epistemic justification to belief.
3. We'd like to stress that we do not take the proposed excusing conditions (DJ and ODJ) as necessary and/or sufficient for blamelessness. Crucially, the problem these cases of radical deception pose for him does not require him to give such conditions. Providing a theoretically motivated account of the intuition of justification in cases of radical deception will do. And Williamson's framework does just this. If those who have pressed the problem of radically deceived agents against the likes of Williamson want to make their objection stick, the onus is back on them to provide reason to believe that the intuition of justification in these cases is not mistaken or at least that Williamson's error theory remains less than fully satisfactory after all.
4. Note that Williamson himself accepts that compliance with the secondary norm entails compliance with the tertiary norm: "Suppose that A complies with DN. Then in any situation, what A does is what some agent who complies with DN does, so A complies with ODN" (forthcoming, p.26).
5. This point against treating radically deceived subjects as merely blameless rather than justified is a familiar one (e.g. Pryor 2001, Bird 2007, Kelp forthcoming-b, Kelp and Ghijzen 2016).
6. Note that the discussion to follow, although mainly focusing on a knowledge norm of belief for simplicity and readability reasons, can in principle be applied to other factive norms of belief as well.
7. This is not to say that they need be action guiding in a strongly internalist sense, that is, what Williamson (2008) dubs perfectly operationalizable norms: norms which are such that one is always in a position to know whether one is complying with them. On the contrary, we take it that most prescriptive norms are not thus operationalizable; take, for instance, the norm: 'Drive 50 km/hour within city bounds!' Whether one complies with this norm is not something that one is always in a position to know. If this isn't immediately obvious, just imagine your speedometer is broken.
8. Note that Geach (1956) argues that there's nothing more to good simpliciter than attributive goodness. For some problems with Geach's line, see Schroeder (2012).
9. This also gets further confirmation in several other places in Williamson (2000, 2014). For instance, according to Williamson, "[k]nowledge sets the standard of appropriateness for belief. [...] Mere believing is a kind of botched knowing. In short, belief aims at knowledge" (2000, p.47). After all, the thesis that knowledge is the evaluative norm of belief, and so a good belief is one that qualifies as knowledge, is an attractive way of making sense of the slogans that mere believing is a kind of botched knowing and that belief aims at knowledge.
10. That said, we admittedly have argued neither that knowledge is indeed the evaluative norm of belief, nor that, in the case of belief, the prescriptive norm serves to ensure that the evaluative norm is likely enough satisfied. We do not mean to deny that these theses are of some importance to knowledge firsters who want to appeal to our account in order to avoid PJD. At the same time, our mission here is not to mount a full defense of knowledge first epistemology in general or some specific version thereof. Rather, our aim is the more modest one of arguing that knowledge first epistemology is not brought down by PJD. The fact that an

independently plausible account of the normativity of belief will allow knowledge firsters to exit PJD, at least modulo certain further assumptions, should be enough to do the trick here.

11. The account is modeled on etiological functionalist accounts, championed e.g. by Ruth Millikan (1984) and in epistemology by Peter Graham (2012).
12. Isn't it more plausible that it is connection to aptness rather than success that explains the positive normative status of competent moves? Perhaps. But note that, for present purposes, this issue is of little consequence. After all, on the account, in the case of inquiry, a move is successful iff apt. In other words, a belief is successful iff it is apt iff it qualifies as knowledge (Kelp forthcoming-a).

References

- Bird, A. (2007). Justified judging. *Philosophy and Phenomenological Research*, 74:81–110.
- Geach, P. (1956). Good and evil. *Analysis*, 17:33–42.
- Goldman, A. (1988). Strong and weak justification. *Philosophical Perspectives*, 2:51–69.
- Graham, P. (2012). Epistemic entitlement. *Noûs*, 46:449–82.
- Kelp, C. (forthcoming-a). Knowledge first virtue epistemology. Carter, A., Gordon, E. and Jarvis, B., editors, *Knowledge First: Approaches in Epistemology and Mind*. Oxford: Oxford University Press.
- . (forthcoming-b). Justified belief: Knowledge-first style. *Philosophy and Phenomenological Research*.
- and Ghijsen, H. (2016). Perceptual justification: Factive reasons and fallible virtues. Mi, C., Slote, M. and Sosa, E., editors, *Moral and Intellectual Virtues in Western and Chinese Philosophy: The Turn Toward Virtue*. New York: Routledge.
- Littlejohn, C. (forthcoming). A plea for epistemic excuses. Dorsch, F. and Dutant, J., editors, *The New Evil Demon: New Essays on Knowledge, Rationality and Justification*. Oxford: Oxford University Press.
- McHugh, C. (2012). The truth norm of belief. *Pacific Philosophical Quarterly*, 93:8–30.
- Millikan, R. (1984). *Language, Thought, and Other Biological Categories*. Cambridge, MA: MIT Press.
- Pryor, J. (2001). Highlights of recent epistemology. *British Journal for the Philosophy of Science*, 52:95–124.
- Schroeder, M. (2012). Value theory. Zalta, E., editor, *The Stanford Encyclopedia of Philosophy* (Summer 2012 Edition). URL = <<http://plato.stanford.edu/archives/sum2012/entries/value-theory/>>.
- Simion, M. (2016). Knowledge-first functionalism. *Unpublished Manuscript*.
- Strawson, P. F. (1962). Freedom and resentment. *Proceedings of the British Academy*, 48:1–25.
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.
- . (2008). Why epistemology cannot be operationalized. Smith, Q., editor, *Epistemology: New Philosophical Essays*. Oxford: Oxford University Press.
- . (2014). Knowledge first. Steup, M., Turri, J. and Sosa, E., editors, *Contemporary Debates in Epistemology* (Second Edition). Oxford: Wiley-Blackwell.
- . (forthcoming). Justification, excuses and sceptical scenarios. Dorsch, F. and Dutant, J., editors, *The New Evil Demon: New Essays on Knowledge, Rationality and Justification*. Oxford: Oxford University Press.

Copyright of Philosophical Issues is the property of Wiley-Blackwell and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.