

ARTICLE

Deep Disagreement, Epistemic Norms, and Epistemic Self-trust

Simon Barker 

Department of Philosophy, Tartu University, Estonia
Email: frigno@gmail.com

(Received 30 September 2022; revised 13 February 2023; accepted 4 March 2023)

Abstract

Sometimes we disagree because of fundamental differences in what we treat as reasons for belief. Such are ‘deep disagreements’. Amongst the questions we might ask about deep disagreement is the epistemic normative one: how ought one to respond to disagreement, when that disagreement is deep. This paper addresses that question. According to the position developed, how one ought to respond to deep disagreement depends upon two things: (i) Whether one remains, in the context of disagreement, permitted to trust oneself in following one’s epistemic norms and in relation to that (ii) whether one has the shared stock of social-epistemic experiences with followers of one’s rival’s norms such that one could reasonably be expected to see their way of doing things as a genuine epistemic alternative to one’s own. Section 1 introduces a model of deep disagreement such that deep disagreements occur when disputants disagree because they follow different epistemic norms. Section 2 surveys the possible theoretical stances to take towards the epistemic-normative question. Section 3 develops an account of what it is to follow an epistemic norm. Section 4 applies this conception of norm-following to the epistemic-normative question. Section 5 discusses two points of clarification.

Keywords: disagreement; deep disagreement; epistemic trust; self-trust; epistemic norms

Introduction

Sometimes we disagree because of fundamental differences in the kinds of consideration we treat as reasons for belief. Such disagreements are ‘deep disagreements’. Amongst the questions we might ask about deep disagreement is the epistemic-normative one, simply put: epistemically speaking, how ought one to respond to disagreement, when that disagreement is deep. This paper addresses that question.

According to the position developed here, how one ought to respond to deep disagreement depends upon two things: (i) Whether one remains, in the context of disagreement, permitted to trust oneself in following one’s norms and in relation to that (ii) whether one has the shared stock of social-epistemic experience with followers of one’s rival’s norms such that it would be reasonable to expect one to see their way of doing things as a genuine epistemic alternative to one’s own.

© The Author(s), 2023. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

Here is the structure of the discussion: section 1 surveys the philosophical questions that deep disagreement poses and introduces a model of deep disagreement such that deep disagreements occur when disputants disagree because they follow different epistemic norms. Section 2 surveys the possible theoretical stances to take towards the epistemic-normative question. Section 3 develops an account of what it is to follow an epistemic norm, whereby ordinarily we come to follow norms through a process of social-epistemic inculcation and internalisation and implement those norms through the manifestation of a form of epistemic self-trust that is grounded in that social-epistemic background. Section 4 applies this conception of norm-following to the epistemic-normative question. Section 5 discusses two points of clarification.

1. Introducing Deep Disagreement

Consider the following case:

CONSPIRACY

Tessa and Connie are colleagues. Connie says to Tessa: ‘Have you read that report by those scientists who bravely broke from orthodoxy. It seems Covid-19 was created in a lab. Tessa replies ‘I saw that, but Covid-19 was clearly not created in a lab. After all, all the other scientists say it wasn’t.’¹

As it happens, Connie spent her childhood in a closeted religious community. She has since left the religion and become deeply engaged with various online conspiracy-theory groups. As a result of this background, Connie has a default preference for explanations of events and phenomena that posit them to be intended and not the result of chance or accident. In believing that Covid-19 was created in a lab, Connie reasons in just this way. In contrast, Tessa was raised in a culturally and intellectually diverse community leaving her without any proclivity to favour one kind of explanation over another. In the absence of any such proclivity, Tessa places greater weight on the testimony of the majority of the scientific community over that of a small group of dissenters.²

Let’s say that disputants are in an ordinary disagreement when they disagree over p , p is a non-normative matter of fact, and that disagreement is best characterised in terms of the disputants having conflicting beliefs about p . Ordinary disagreements are just that – the stuff of everyday intellectual life, social-epistemic interactions, conversations, debates, and conflicts. On the surface, CONSPIRACY resembles an ordinary disagreement. Tessa and Connie disagree over the proposition ‘Covid-19 was created in a lab’, that proposition refers to non-normative matters of fact, and it is apt to describe Tessa’s and Connie’s attitudes towards this proposition in terms of belief. Whilst, however, Tessa and Connie’s disagreement resembles an ordinary disagreement at the level of belief, there is an underlying difference in terms of the kinds of consideration that

¹At the time of writing, the prevailing view in the scientific community is that the Covid-19 virus has natural origins. The US Department of Energy and FBI have recently come out as having, respectively, low and moderate confidence that the Covid-19 pandemic started from a lab leak. Even these institutions, however, have not made the further claim that the virus itself was created in a lab – a claim that remains the province of conspiracy theorists. (Thanks to an anonymous reviewer for pointing out the shift in status of this debate.)

²For empirical work on the connection between belief in conspiracy theorising and so-called ‘intentionality bias’, see Brotherton and French (2015).

each treats as relevant to determining the truth. I.e., Connie has a predilection for believing intentional explanations of events and phenomena, Tessa does not share this predilection. Crucially it is this underlying difference that explains how the two have come to disagree.

CONSPIRACY is an example of *deep disagreement*. Other examples discussed in the literature include disagreements between practitioners of alternative and conventional medicine; disagreements over climate change (Kappel 2021); disagreements between those who follow scientific principles and those who follow the authority of religious texts (Lynch 2010; Kappel 2012; Matheson 2021; Pritchard 2021); and disagreements between people on opposite sides of the political spectrum (Pritchard 2021).

Discussion of such disagreements has coalesced around several questions. We might describe these as follows:

The Incommensurability question:

Are deep disagreements inherently resistant to rational resolution? (See Fogelin 1985; Lynch 2010; Matheson 2021; Pritchard 2021; Ranalli 2020.)

The Meta-epistemological question:

If deep disagreements are rationally irresolvable, does their possibility force us to accept some form of epistemic relativism? (See MacIntyre, 1977; Lynch 2010.)

The Practical question:

How, practically speaking, ought we to respond to deep disagreements so as to avoid unconstructive impasses in contexts where coordination between different people and groups is required? (See MacIntyre 1977; Lynch 2010; Kappel 2012.)

The (Epistemic) Normative question:

How, epistemically speaking, ought one to respond to the realisation of deep disagreement? (See Matheson 2021; Ranalli 2020; Kappel 2021.)³

It should not be surprising that these questions are, at times, treated as interconnected. For instance, Matheson (2021) and Ranalli (2020) address the incommensurability question via consideration of whether those party to deep disagreement are permitted to continue believing the relevant propositions or whether they ought all to suspend judgement on those propositions. Both, then, proceed from the presumption that the normative and incommensurability questions are interdependent. MacIntyre (1977), Lynch (2010) and Kappel (2012), on the other hand, introduce the practical question to discussion just because they posit that deep disagreements are irresolvable from the perspective of epistemic rationality. And so on.

At the same time, we should recognise that answering any one question does not necessarily settle any of the others. For instance, epistemic relativism need not entail that an individual agent is permitted to hold conflicting beliefs. So, the normative question remains open, even if the answer to the meta-epistemological question is yes. Deep disagreements may be rationally irresolvable without that implying anything specific about how one should respond to finding oneself in deep disagreement – practically or epistemically. And so on.

This paper addresses the normative question. Before turning to discuss that question proper, it will help to have an appropriate characterisation of deep disagreement in

³Clearly, the practical question is also a normative question. Since this paper is concerned with questions about the epistemic significance of deep disagreement, however, I restrict talk of normativity solely to epistemic normativity.

hand. I suggest that there are two features such a characterisation needs to capture if it is to be fit for that purpose. These can be described as follows:

No Theory Needed (*NTN*): One can be in in deep disagreement without realising one is in deep disagreement or being in a position to realise one is in deep disagreement.

Normatively Distinct (*ND*): The loci of deep disagreements are not beliefs upon non-normative matters of fact, but some general feature of our belief-forming practices relevant to the epistemic status of beliefs about non-normative matters of fact.

NTN is important because many of us lack the theoretical and conceptual resources – let alone time and opportunity – required to uncover the kind of systematic difference in epistemic practices that we might associate with *CONSPIRACY* or the other examples from the literature. Thus, if it were required that disputants realise or be able to realise that their disagreement is deep, vanishingly few disagreements would count as deep – including those taken to be paradigmatic of the class. *ND* is important on the presumption that those party to deep disagreement may face normative issues distinct from those posed by ordinary disagreement. Since, then, the loci of ordinary disagreements are beliefs about non-normative matters of fact, the loci of deep disagreements must lie with some other epistemically significant feature of the practices by which those party to such cases come to disagree.

With these features of deep disagreement in mind, what I suggest is that a natural way of describing deep disagreement is in terms of the *epistemic norms* that disputants followed in coming to disagree at the level of belief.⁴

In general, following Boghossian (2008) and Littlejohn and Turri (2014), we might gloss epistemic norms as rule-like statements of the following form:

Where F designates some epistemic act such as believing, inferring, trusting, and so on,

Under condition C, one ought to/ought not to/is permitted to F

Or imperatives of the form:

If C, F!/don't F!

Let's say too that, when genuine epistemic norms have it that one ought to/ought not to/is permitted to F – the mandates of those norms will be indexed to whatever are the intrinsic success conditions of the relevant act. Most fundamentally, for instance, since belief aims at truth, genuine norms of belief encode conditions under which one ought to/ought not to/is permitted to believe p – *for the sake of believing truths and not falsehoods*.

In the simplest sense, I suggest, to say someone follows some specific norm is to say that they somehow regulate the relevant epistemic actions in accordance with what they

⁴The literature offers two other models: 'The Wittgensteinian model', whereby deep disagreements arise because disputants have conflicting *hinge-commitments* (see Fogelin 1985; Pritchard 2021; Ranalli 2020); and The 'Principle model', whereby deep disagreements are disagreements over conflicting *epistemic principles* (see Lynch 2010; Kappel 2012; Lynch and Silva 2016; Matheson 2021). For reasons of space, this paper does not discuss the merits and demerits of these two models (for a thorough discussion along those lines, see Ranalli 2021). However, the epistemic-norms model offered here can be seen as a kind of synthesis of the Wittgensteinian and Principle models.

perceive to be the prescriptions of that norm. Thus, to say a person perceives the situation as one in which a norm they follow *requires* or *permits* them to believe-*p* is to say that they perceive the situation as one in which they have reason to believe-*p*. To say they perceive the situation as one in which that norm *proscribes* against believing-*p* is to say they perceive the situation as one in which they do not have reason to believe-*p*. I shall give a fuller account of what it means to follow an epistemic norm in section 3. For now, the important points are that following an epistemic norm in the sense described is (i) something an agent does, not necessarily a matter of what they believe, and (ii) constitutively connected to what an agent treats as epistemic reasons. The former I take to be self-explanatory. The latter follows from the idea that questions about what we ought to believe are questions about what we have reason to believe. With this much said, we can characterise deep disagreement as follows:

Deep Disagreement

Where *S* follows some epistemic norm *N* and *R* follows some epistemic norm *M*, then:

If *S* believes *p* and *R* believes $\sim p$, AND that disagreement is explained by *S*'s competently following *N* and *R*'s competently following *M*, then, the disagreement between *S* and *R* is a **deep disagreement**.⁵

Since following a norm is something that one does and not necessarily what one believes, and deep disagreements so defined come about via disputants following different norms, this definition fits with *NTN*. Since the loci of deep disagreements on this definition are differences in norms and since norm-following is directly relevant to questions about what we ought to believe, this definition fits with *ND*. To illustrate, let's return to *CONSPIRACY*. Given Connie's favouring of intentional explanations, we might say that Connie follows the norm:

Connie-norm: If *p* and *q* cannot both be true, *p* is an intentional explanation of *E*, *q* is a non-intentional explanation of *E*, *E* is unlikely to be true unless *p* or *q* are true, and *E*, then, one ought to believe *p* (and not *q*).

Given that Tessa does not, by default, favour any particular kind of explanation, we might say that Tessa follows the norm:

Tessa-norm: If *p* and *q* cannot both be true, *E* is unlikely to be true unless *p* or *q* are true, there is independent evidence that *p* is more likely true than *q*, and *E*, then, one ought to believe *p* (and not *q*).

Crucially, as the case goes, Connie's norm says that a follower of that norm ought to believe 'Covid-19 was created in a lab'. Tessa's norm mandates a follower of that norm to believe 'Covid-19 was not created in a lab'. What explains the disagreement, then, is nothing so ordinary as a difference in evidence or personal-performance error.

⁵An anonymous reviewer suggests we might expand this definition in two ways, (i) to allow for disagreements that involve differences in credence (ii) to allow for disagreements where one party believes *p* and the other suspends judgement on *p*. In principle I see nothing wrong with expanding the account of deep disagreements developed in this paper in these ways. However, for the sake of simplicity, I shall continue to employ the definition in the main body of the text.

Instead, the disagreement is explained by the fact that Connie and Tessa *competently followed the prescriptions of rival norms*. Thus, even if neither Connie, nor Tessa have the theoretical resources that would allow them to think of it as such, their disagreement counts as a deep disagreement. Additionally, since deep disagreements on this definition entail – not only a difference at the level of norms – but a conflict in the beliefs that the disputants come to hold via following those norms, deep disagreements so understood are genuine disagreements. And since they are genuine disagreements, deep disagreements raise the normative question – i.e., epistemically speaking, how ought one to respond to the realisation of disagreement, when that disagreement is deep?

With this characterisation in hand, we are equipped to consider the normative question proper. Let's start by considering the general positions that theoreticians might take on that question.

2. Rational Options and Theoretical Positions

Since deep disagreements are constituted by differences at two levels (i.e., the level of norms and the level of beliefs), the ways in which disputants might respond to deep disagreement are complicated in a way they are not in ordinary disagreement. In the case of ordinary disagreement, the normative question is whether disputants ought to revise or retain their beliefs about the disputed propositions. In the case of deep disagreement, there are two such questions: Should the disputants revise or retain their original beliefs? And should the disputants continue to follow their respective norms? In contrast to the two possible responses to ordinary disagreement, then, there are four possibilities for how disputants might respond to deep disagreement:

- (a) Stick by norms, retain beliefs.
- (b) Stick by norms, revise beliefs.
- (c) Reject norms, retain beliefs.
- (d) Reject norms, revise beliefs.

As far as rational options go, we can immediately rule out (c). In the absence of support for believing as they do via some other uncontested norm, it will not be rational for someone engaged in a deep disagreement to reject their norms yet retain the beliefs formed in accordance with those norms. Consider Tessa in *CONSPIRACY*. As the case is described, any reason that Tessa has for believing 'Covid-19 was not created in a lab' comes by way of her being permitted to follow the norm by which she came to hold that belief. Tessa's being rationally permitted to hold that belief, thus, entails that she is permitted to follow the relevant norm. Correspondingly, if, in the face of deep disagreement, Tessa ought to reject the contested norm, then, she will lose whatever reason she had to believe 'Covid-19 was not created in a lab' that came by way of her following that norm. And so, since she has no support for believing 'Covid-19 was not created in a lab' from another uncontested norm, it cannot be the case that Tessa ought to reject her norms and retain that belief. In other words, other things being equal, (c) is not a rationally viable response to the disagreement. More generally, this helps us to recognise the following principle:

Total Revision: *Ceteris paribus*, in light of deep disagreement, if S ought to reject their epistemic norm N, then, for any disputed belief p that S came to hold via N, S ought to revise their belief p.

(The *ceteris paribus* clause here reflects the possibility alluded to above that one may have an epistemically appropriate route to the same belief via some other norm.)

I shall return to the significance of this principle shortly. First, a note on the plausibility of (b). Like (c), (b) calls for a different response at the levels of norms and beliefs (i.e., stick by norms, revise beliefs). However, unlike (c), (b) will be appropriate in some cases of deep disagreement. This is for the simple reason that epistemic norms can be defeasible. To say that norms are defeasible, is just to say that one can be in the situation where one has no reason to reject one's norms in any general fashion, but one ought to withhold from following the prescriptions of those norms in the specific context. The mixed response at (b) should be understood in this way. (To pre-empt the later discussion, (b) may be appropriate in *some* cases where the disputants are not in a position to realise that their disagreement is deep.)⁶

Given these points, then, (a), (d), and (b) – properly understood – are all theoretically viable responses. Leaning on terminology from the literature on ordinary disagreement, let's label (a) the *steadfast* response to deep disagreement, and (d) the *conciliatory* response.⁷ And, in so far as it requires conciliation at one level and steadfastness at the other, let's label (b) the *mixed* response to deep disagreement. With that said, we can categorise the possible views of deep disagreement in terms of the range of cases in which they prescribe or permit these different responses. So, we can call accounts that prescribe the conciliatory response in all cases *uniform-conciliationist* views of deep disagreement. Accounts that permit at least one disputant to stay steadfast in all cases we can call *uniform-steadfast* views. (Or, at least in all cases where it is not so that neither disputant was permitted to follow the relevant norms to start with.) Finally, we can recognise the possibility of *non-uniform* accounts of deep disagreement, whereby which of the steadfast, conciliatory, or mixed responses is appropriate will depend upon the circumstances of the particular disagreement. In sections 3–5, I offer an account of what it is to follow a norm that supports a non-uniform view of deep disagreement.

Before doing so, it is worth considering some motivations we might have for endorsing a non-uniform account of deep disagreement independently of the merits of the account to follow.

Consider the principle of Total Revision. On a uniform-conciliationist account of deep disagreement, this principle would apply in all cases of deep disagreement. In a case such as CONSPIRACY, however, this would have wide-ranging and comprehensive sceptical consequences. Why? Because, whilst the disagreement between Tessa and Connie concerns biological and political matters, the norms that explain that disagreement are not specific to those domains. As described, Connie's norm mandates favouring intentional explanations in any and all instances where such explanations are salient by her lights. Tessa's norm rules out such a general bias. Thus, if uniform-conciliationism is correct, Tessa and Connie should, according to Total Revision, withhold judgement on any and all issues in which non-intentional and intentional explanations support different conclusions. Just as Tessa's and Connie's norms are fully general, however, they are such that we might expect swathes of disagreement between

⁶An anonymous reviewer points out that (b) may also be appropriate if one's norms includes some kind of disagreement-norm that mandates revising one's belief in case of deep disagreement. In such a situation, one may be in a position such that one should both stick with the norm and revise one's belief. In this case, (b) may be appropriate even if some norms are not defeasible.

⁷As far as I can tell, the earliest published use of this distinction is from Christensen (2009).

followers of those norms across many – if not all – domains. And so, we might expect too that a conciliatory response to the disagreement would require the pair to revise a great many of their beliefs.

Of course, as per the *ceteris paribus* clause on Total Revision, these sceptical consequences may be mitigated if Tessa and Connie have support for any of the relevant beliefs that is grounded in some other undisputed norm. Here, though, it is worth reiterating that the norms in question are both general and fundamental (just as is the practice of inference to the best explanation that they in some sense specify). And, I would suggest, the more general and fundamental a norm is, the less likely it is that any belief formed via following that norm will be well-supported by some other (undisputed) norm. If that is correct, then, whilst some of the beliefs Tessa and Connie have formed via their rival norms may be safe from the sceptical consequences of Total Revision, the majority will not. And, again, since we would expect there to be swathes of disagreement between followers of the two norms, we would still expect the sceptical consequences of a conciliatory response to be great.

I make these points here since much of the discussion in the literature on deep disagreement has focused on cases such as CONSPIRACY. As we saw in section 1, common examples include disagreements between practitioners of alternative and conventional medicine, those who follow scientific principles of inquiry and those who follow the authority of religious texts, people on opposite sides of the conservative/progressive political divide, and so on. All these, though, are examples in which we might expect the rival norms to be general in scope and cover a range of issues broad enough that any sceptical consequences of the disagreement will be comprehensive – if not global.⁸ An additional feature that is common to these cases, including CONSPIRACY – and one that I think stems from just the kinds of sceptical worries at issue – is that intuitively *at least* one of those party to such cases ought to stay steadfast both in norms and in beliefs.⁹ Indeed, amongst those who discuss the normative implications of deep disagreement, steadfastness is the more popular choice. So, for instance, Ranalli writes:

it's not clear at all that in cases of deep disagreement so understood you get higher-order evidence that you've made a mistake, since it's hard to see how you could have good enough evidence for thinking that your opponent is your peer. So, it seems as if the rational response to the disagreement is for them to retain their attitudes of trust and thus continue to disagree. (Ranalli 2020: 5000)

Kappel writes:

In general terms the HOE account says the following about deep disagreement. In the general case, disagreement is *prima facie* higher order evidence of performance errors or principle errors, and in response to this one should modify one's credence towards less certainty. However, in [transparent] deep disagreement, this

⁸If the reader is not convinced that the sceptical consequences of conciliation in one-off cases are enough to challenge uniform-conciliation, it is worth remembering that the view says that conciliation is the proper response to *all* deep disagreements. Presuming one could encounter numerous extreme deep disagreements such as CONSPIRACY, then, uniform-conciliationism would still seem to cede the mountain to the sceptic.

⁹I take it the pro-steadfast intuition is stronger for Tessa than for Connie – just because, plausibly, Connie's norm was defective to start with. Presuming, then, that the realisation of disagreement cannot improve one's epistemic standing, Connie won't be rationally permitted to stay steadfast. For this reason I focus upon Tessa in the rest of the paper.

higher order evidence is undercut by evidence stemming from knowledge of the nature of the disagreement. In so far as deep disagreements are symmetrical, neither party should conciliate. (Kappel 2021: 1045)¹⁰

It is important to reiterate that Ranalli and Kappel (and Matheson) are focusing on cases such as CONSPIRACY where the disagreement involves norms that are general in scope. By focusing primarily on such cases, however, there is a danger that we are too easily tempted into endorsing, not just the steadfast response in the particular case, but a uniform-steadfast account of deep disagreement. Too readily giving into that temptation, I think, would be a mistake. For, the differences between those party to deep disagreement are not always so extreme as in CONSPIRACY or the various other examples listed. Indeed, it strikes me that some of the most interesting examples of deep disagreement will be ones in which the disputants are in many ways, so to speak, following the same rulebook, yet are embroiled in deep disagreement nonetheless. In these cases, it is far less clear that either side to the dispute should stay steadfast. Nonetheless, such disagreements can be just as intractable and resolving them just as problematic as the cases typically discussed. To illustrate these thoughts, consider the following case:

DOCTORS

Doctors Empany and Evitt are treating a patient with condition-X. Both are highly trained and experienced. Both have carefully assessed the clinical evidence and discussed the treatment options with the patient. Having consulted the relevant medical guidelines, Empany concludes that treatment-A is best, Evitt that treatment-B is best.

Empany and Evitt based their conflicting beliefs on the same evidence but differed in their underlying approaches to that evidence. Both thought the clinical evidence best fits the conditions under which the medical guidelines recommend treatment-B. However, in discussion, the patient expressed considerable concern about treatment-B. In coming to her conclusion, Empany placed weight upon these expressions of concern. It was this evidence-from-empathy that led her to conclude that treatment-A would be better. In contrast, Evitt put aside his empathy with the patient in favour of strictly following the medical guidelines.

As it happens, the college Evitt attended teaches students to base medical beliefs only upon high-quality scientific and clinical evidence. The college Empany trained at teaches students to base medical beliefs primarily upon consideration of high-quality scientific and clinical evidence, but to also give weight to evidence-from-empathy. Crucially, doing so sometimes, though not often, leads Empany to forming beliefs that run counter to the recommendations of the relevant guidelines. This is one such case.¹¹

¹⁰Though less explicitly framed as such, Matheson (2021) can also be understood as defending steadfastness when he suggests that any disputant who follows *true* epistemic principles can stick with those principles in the face of deep disagreement.

¹¹This case is a simplification of differences between models of medical practice that emphasise a role for empathy – including narrative medicine and the biopsychosocial model – and those that exclude such sources of evidence – including ‘evidence-based medicine’ (EBM) and the biomedical model. For discussion of the differences between EBM and narrative medicine, see Solomon (2015); for discussion of differences between the biomedical and biopsychosocial models, see Gifford (2017).

In *DOCTORS*, Empany and Evitt disagree about the truth of a narrow set of propositions, i.e., ‘treatment-A is best’ and ‘treatment-B is best’. However, the disagreement is not explained by their having different evidence, referring to different guidelines, or by any personal performance error when assessing the relevant evidence. Instead, the roots of the disagreement lie in the weight each gave to their empathy with the patient’s expressions of concern about treatment-B. What is more, in doing so, each followed the general understanding of how medical beliefs ought to be formed into which they had been inculcated during their medical training. In this respect, I would suggest, the difference in how the pair came to their beliefs can be characterised in terms of epistemic norms.

So, we might say, Empany follows the norm:

Empany-norm: If, having considered the available clinical and scientific evidence as well as evidence gained from empathising with the patient, your clinical judgement is that *p* is likely to be true, you ought to believe *p*; if not, not.

Evitt follows the norm:

Evitt-norm: If, having considered the available clinical and scientific evidence, your clinical judgement is that *p* is likely to be true, you ought to believe *p*; if not, not.

Since it is the difference in these norms that explains the disagreement, similarly to *CONSPIRACY*, *DOCTORS* counts as a deep disagreement as defined. Unlike *CONSPIRACY*, however, Empany’s and Evitt’s norms do not prescribe how one ought to or ought not to form beliefs in all circumstances – but only when deliberating on the truth of propositions that fall within the medical domain. In this respect, we might say that this is a *domain specific* deep disagreement. Reflecting that, to leverage a concept from the wider literature, we might note also that the pair should think of each other as *epistemic peers* – at least that is independently of the difference in their norms. Where, as it is most commonly described, disputants are peers when there is rough symmetry in (i) their access to and familiarity with the relevant evidence and (ii) factors that affect how competent they are at processing that evidence. These include, for instance, the possession of relevant virtues and cognitive abilities and circumstantial factors that might impair or enhance disputants’ ability to process the evidence in the particular context (see, for instance: Christensen 2009; Kelly 2010; Lackey 2010; Matheson 2015).

Why think that Empany and Evitt should think of each other as peers in this way? For one thing, Empany and Evitt have equal access to and familiarity with both the clinical and scientific evidence *and* the information gained by empathising with their patient. For another, there are no circumstantial factors that might enhance or impair either of their ability to process this evidence. Moreover, both are aware of these features of the case. Thus, the question of whether the pair should think of each other as peers is the question of whether they have evidence that they are similarly competent in a more general sense. There are two reasons for thinking they will have such evidence – moreover there are two reasons for thinking that each will have such evidence *by the light of their own norms*:

Firstly, though allowing space for evidence-from-empathy, Empany’s norm still mandates that one place significant weight on the clinical and scientific evidence. Secondly, and more significantly, for Empany to competently follow her norm, she

will have had to have honed her judgement of the import of evidence-from-empathy through medical training and practice. More specifically, she will have honed it – even if not explicitly – against standards and norms operative in the medical community. These will include the standards explicitly formulated in medical guidelines, as well as those explicitly or implicitly instantiated in the practice of other qualified medical professionals, including someone such as Evitt. Given this, though, one would not expect a competent follower of Empany’s norm to disagree widely with a competent follower of Evitt’s norm. If that is so, however, then given their similar experience and competence, we would also expect Empany and Evitt’s track records to provide considerable evidence that both of the pair are competent *by the lights of either Empany’s or Evitt’s norm*. Presuming that the pair have some access to such evidence, then, and in the absence of some other difference-making feature, the pair should think of each other as epistemic peers in the way described.

What is the significance of the differences between the two cases?

For one thing, a conciliatory response to a domain specific deep disagreement – even granting the principle of Total Revision – would not have as comprehensive sceptical consequences as in a case such as CONSPIRACY. The anti-sceptical motivation for a uniform steadfast view of deep disagreement stems from the thought that, if all cases of deep disagreement are like CONSPIRACY, then, given the sceptical consequences of conciliating in such cases, allowing that all deep disagreements require a conciliatory response would be to cede the mountain to the sceptic. However, the possibility of domain-specific deep disagreements undercuts this anti-sceptical motivation. Granted the consequences of conciliating for, say, Empany and Evitt would still be extensive – but they would not be so-near global. Thus, we might allow the possibility that the conciliatory response is appropriate in some cases without such concessions to the sceptic as would be were conciliation always the appropriate response.

For another thing, the possibility that the relevant norms of those party to a deep disagreement might be sufficiently close that they should view each other as peers undercuts the intuition fuelled by cases such as CONSPIRACY that at least one party to the disagreement ought not to harbour any doubts about their practices. In cases such as CONSPIRACY, it is appealing to think that someone in Tessa’s position should not doubt her practices for the reason that, by her lights, Connie is unlikely to seem in any way a competent agent. But such intuitions do not hold up in a case such as DOCTORS. As noted, in this case, despite their conducting their medical practices according to conflicting epistemic norms, there is reason to think that each party will have significant evidence that the other is in fact competent. But that being so, the intuition that, say, Empany ought not to question herself in respect to the practices by which she came to believe ‘treatment-A is best’, simply because Evitt follows different norms, has far less traction than the parallel intuition re Tessa in CONSPIRACY.

In light of these differences, then, we are faced with the following question: Is it possible that the appropriate response for Tessa in CONSPIRACY is to be steadfast, whilst Empany and Evitt in DOCTORS ought to be conciliatory? We might refer to this question as the ‘Normative Problem of Deep Disagreement’. To resolve it, I suggest, we need to look to the defining characteristic of deep disagreements: i.e., that they are explained by the disputants following rival epistemic norms. More specifically, we need to consider (i) how we come to have the norms we do and (ii) how we implement those norms once we have them. Or, to put it more simply, we need to consider the question of what it means to *follow* an epistemic norm. I turn to this question in the next section.

Before doing so, it is worth considering whether we already have an explanation of how the appropriate response to disagreement for Empany and Evitt could differ from the appropriate response for Tessa. That is, the norms in *DOCTORS* are domain-specific, whilst those in *CONSPIRACY* are not. And, connected to that, Empany and Evitt should think of each other as epistemic peers, whereas Tessa is under no such pressure to think of Connie as her peer. This seems correct.¹² As it stands, though, this difference in the two cases offers only a *partial explanation* of why the appropriate response to the disagreement for Empany and Evitt will differ from that for Tessa. For one thing, it remains an issue of debate within the literature as to whether *conciliationism*, *steadfastness*, or some *non-uniform* view is the correct account of peer disagreement.¹³ And, if questions about the significance of peer disagreement more generally remain unsettled, we cannot *only* appeal to the fact that Empany and Evitt are peers to explain how they should respond to their disagreement – one way or the other. For another, the discussion of peer disagreement has largely focused on cases of *ordinary* disagreement between individuals. That being so, however, then even if the question of peer disagreement were settled in ordinary cases, we cannot straightforwardly presume that whatever is the correct account of ordinary disagreement will apply to cases of deep disagreement. If, then, we want a *full explanation* of why the appropriate response for Empany and Evitt should differ from the appropriate response for Tessa to hers, we want also an account of *why* and *how* the difference between the two cases re attributions of peerhood is significant in the specific context of deep disagreement. It is in the service of providing just such an explanation that I suggest we need to consider the question of what it means to follow an epistemic norm. So, let's do that.¹⁴

3. Epistemic Norms as Social Practices

In the first half of the paper, I suggested that we can usefully characterise deep disagreements as disagreements that occur when disputants have come to hold conflicting beliefs via their following different epistemic norms. In relation to that, I suggested that whatever is the correct response to any case of deep disagreement will turn

¹²An anonymous reviewer points out another feature of the cases that may be salient. What we know about the number of people (including peers) on each side of a disagreement can make a difference to how we ought to respond to that disagreement. But, where we might suppose the numbers play out relatively evenly in a case such as *DOCTORS*, it seems that those who follow conspiratorial norms like Connie's will be well outnumbered by those who don't. This difference between the cases, then, may explain the intuition that Tessa may stay steadfast whilst Empany and Evitt ought to be conciliatory. I think there is some pull to this suggestion. As the reviewer notes, however, this difference is more an accident of the examples than a feature of deep disagreements in general. We can find cases similar in scope to *CONSPIRACY* that do not have this number difference. Since, though, I think *CONSPIRACY* is in other ways a good illustration of the extreme deep disagreements we are interested in, I shall continue with this case. If, the reader is unhappy with this, I ask them to imagine a preferred case, similar in scope but without this feature.

¹³Notable defences of conciliationism include Elga (2007), Christensen (2011), and Matheson (2015). Defenders of steadfastness include Enoch (2010) and Schafer (2015). Non-uniform views have been forwarded by Feldman (2009), Kelly (2010), Lackey (2010), and Faulkner (2016).

¹⁴An anonymous reviewer suggests that we could alternatively frame the discussion around the question: 'Why should Empany/Evitt think of each other as peers, whilst Tessa need not think of Connie as her peer?' If the reader prefers, I suggest that the discussion could also be read this way; i.e. the reasons given as to why Tessa is permitted to remain steadfast whilst Empany/Evitt should be conciliatory can also be thought of as reasons why Tessa need not think of Connie as her peer whilst Empany/Evitt should think of each other as such.

upon what it means to follow an epistemic norm. What I want to suggest in this section is that to follow an epistemic norm is to somehow regulate one's beliefs in accordance with one's norms. Where, for most of us at least, the regulation of beliefs in accordance with epistemic norms can be explained by one's having been inculcated into normatively loaded ways of *perceiving* the world via exposure to the epistemic practices of the communities, and members of the communities, of which we are a part. And, once so inculcated, *trusting* oneself in respect to the reliability of those ways of perceiving the world. We might call this the 'inculturation and self-trust' model of norm following. As we shall see later, what it means to have the kind of self-trust described is directly relevant to the normative problem of deep disagreement.

Precedent for this 'inculturation and trust' model can be found in a wide range of literature. This includes McDowell's (1978, 1979) work on moral virtue, Fricker's (2007) work on testimony and epistemic injustice, accounts of trust and self-trust from Faulkner (2011) and Jones (2012a, 2012b), and Sripada and Stich's (2012) work on the psychology of social and moral norms. Given the breadth of this literature, I shall not endeavour to lay out a systematic account of this understanding of norms and norm acquisition. Rather, I draw on the aforementioned work to explore the features of norms and norm-following that I take to have significance when it comes to discussion of the normative significance of deep disagreement.

3.1. *The inculturation and self-trust model of norm-following*

Discussing how individuals acquire norms in a general sense, Sripada and Stich write that:

There is excellent evidence indicating that norms exhibit a reliable pattern of *ontogenesis*. Regardless of their biological heritage, almost everyone acquires the norms that prevail in the local cultural group in a highly reliable way. In no human group is it the case that some individuals reliably acquire the prevailing norms while many others don't. (Sripada and Stich 2012: 289)

So too, I suggest, we should expect this kind of social-sourcing to happen with epistemic norms. Ordinarily, individuals will not come to follow the norms they do entirely on their own resources and in opposition to the community to which they belong. Rather, the typical route will be one of inculturation via the individual's exposure to the way others in their community conduct their epistemic practices. As the agent observes and interacts with others in the community to which they belong, they witness how those others attribute authority, who they trust and don't trust, under which conditions people form beliefs, make assertions, offer testimony, and so on. Immersed in this stream of experience, and later becoming an active participant in it, the agent comes to see these practices as *the way things are done* and, so, develops a sensibility about and sensitivity to the epistemically salient features of different situations. Fricker offers a useful description of what I have in mind here in her account of how we develop a (virtuous) testimonial sensibility:

we should think of the virtuous hearer's sensibility as formed by way of participation in, and observation of, practices of testimonial exchange. There is, in the first instance, a passive social inheritance, and then a sometimes-passive-sometimes-active individual input from the hearer's own experience. Together the individual and

collective streams of input are what explain how our normal unreflective reception of what people tell us is conditioned by a great range of collateral experience – our informal background ‘theory’ of socially situated trustworthiness. (Fricker 2007: 82)

A crucial element of this process is that the agent witnesses not only the patterns by which individuals conduct their own epistemic affairs, but also the responses to those activities by others – especially the affective dimensions of those responses. Just as significant, for instance, as seeing who and when individuals trust, is seeing the association of reactive attitudes such as resentment and betrayal with failures to make good on that trust, and gratitude when it is acted on.¹⁵ Similarly for other epistemic actions. If an individual consistently fails to form beliefs or make assertions that reflect the available evidence (or whatever else is considered epistemically salient in the community in question) they are likely, not only to be corrected, but to be subject to criticism. The further their actions stray from the norms, or the greater the practical consequences of their doing so, the harsher and more censorious this criticism is likely to become.¹⁶ These experiences, I suggest, provide the catalyst for the development of the agent’s sense of the way things are into a genuinely normative and genuinely epistemic sensibility. Thus, as one comes to inhabit these communal practices of praise and blame, so too one comes to see compliance with the associated patterns of conduct, not just as something that people do, but something that people, including oneself, *ought* to do. Sripada and Stich describe this in the moral context in terms of the internalisation of norms – where internalisation is a term from anthropology and sociology to describe ‘a characteristic style of motivation in which the individual values compliance with moral rules even when there is no possibility of sanction from an external source’ (Sripada and Stich 2012: 292). As Sripada and Stich point out, the internalisation hypothesis explains why, once indoctrinated into a rule, ‘people exhibit a lifelong pattern of highly reliable compliance with the rule’ (Sripada and Stich 2012: 292).

This gives us an idea of how we come to follow the norms we do. The second question is how we implement those norms once acquired. Echoing aspects of McDowell’s account of the moral virtues and Fricker’s account of testimony, I suggest that we can think of this in terms of the agent having acquired a normatively loaded way of *perceiving* the situations they find themselves in. So, McDowell writes:

In moral upbringing what one learns is not to behave in conformity with rules of conduct, but to see situations in a special light, as constituting reasons for acting; this perceptual capacity, once acquired, can be exercised in complex novel circumstances, not necessarily capable of being foreseen and legislated for by a codifier of the conduct required by virtue. (McDowell 1998: 85)

And Fricker:

The virtuous hearer does not arrive at her credibility judgement by applying pre-set principles of any kind, for there are none precise or comprehensive enough to

¹⁵See Holton (1994), Jones (1996), and Faulkner (2011) on the reactive attitudes and interpersonal trust.

¹⁶Sripada and Stich (2012: 291–5) provide a useful overview of findings from anthropology and economics supporting the idea that the internalisation of norms includes also the development of intrinsic motivations to *punish* norm-violations.

do the job. She ‘just sees’ her interlocutor in a certain light, and responds to his word accordingly. (Fricker 2007: 75)

On this understanding applied to epistemic norms, just as we might come to subscribe to certain norms without explicitly reasoning upon their content, beliefs formed via norm-following do not necessarily depend upon inference from intentions and beliefs about one’s norms and the situations one finds oneself in. The perceptual analogy helps us to see how. For, typically at least, the beliefs we form via perception do not involve reasoning from beliefs about the reliability of our perceptual faculties and experiences. Likewise, on this analogy, the beliefs we form via following our epistemic norms, typically at least, come about through, to use Fricker’s terms, ‘spontaneous and unreflective’ responses to consideration of the circumstances in which we find ourselves (2007: 80). So, for instance, considering the question ‘what is in front of me?’ I turn my attention to the object in front of me. I have the visual experience of a grey table, *in so doing*, I come to believe ‘there is a grey table in front of me’. Considering the question ‘how many books are on that table?’, I see that there are two books on the right side, two on the left, and none in the middle, *in so doing* I come to believe that ‘there are four books on the table in front of me’. In the second instance, my belief is formed, in part, because I have internalised epistemic norms (and rules of mathematics). However, I need not consciously consider those norms and rules when I form the belief. Just as I don’t consciously reflect upon the reliability of my perceptions when I form the belief in the first instance. Instead, we might say that in considering the question ‘how many books are on the table’ I have primed myself to pay attention to those features of the situation that correspond to the application conditions of the relevant norms. When I *see* those application conditions to apply, I respond as those norms prescribe. This way of perceiving things is so ingrained into my way of thinking trusting myself in respect to it is automatic.

This description of how we manifest our internalisation of epistemic norms has echoes of affective conceptions of epistemic self-trust and interpersonal-trust offered by, respectively, Jones and Faulkner. So, Jones writes:

those with intellectual self-(dis)trust, in a domain, experience an emotionally laden perception of the situation. ... This perception typically has a distinct phenomenology and sets in train the cognitive functional roles characteristic of affect: *to control salience, direct thought, and influence inference*. (Jones 2012b: 5, italics added)

And Faulkner:

where the norms of trust are internalized, the subject’s perception of a situation wherein the norms apply ... will be structured by the prescriptions of the norms. ... That is, *the subject’s perception of a situation wherein the norms apply will contain within it a judgement about what ought to be done*. (Faulkner 2011: 185, italics added)

In light of these parallels, then, I would suggest that it is apt to describe the way in which we (typically) come to follow and implement epistemic norms as the manifestation of a certain kind of *socially grounded epistemic self-trust*. It is *socially grounded* in the way that one comes to follow the relevant norms via inculturation into and internalisation of socially sanctioned ways of regulating one’s epistemic actions. It is *epistemic*

self-trust in the way that one unreflectively depends upon one's own perceptions of the relevant features of the world to reliably track the application conditions and prescriptions of the relevant norms.¹⁷

Self-trust of this kind, I propose, will be *well-founded* when three conditions are met:

1. The norms in question are epistemically adequate.
2. One's perceptions reliably track the application conditions and prescriptions of the relevant epistemic norms.
3. The reliability of one's perceptions is grounded in one's inculturation into and internalisation of the relevant epistemic norms.

The first condition is important in so far as the normative status of our beliefs (and other epistemic actions) will depend in some way upon the content of the norms that we follow. Since self-trust of the kind described is closely tied to believing, then, it seems reasonable to suppose that the well-foundedness of one's self-trust will be closely tied to the content – and adequacy – of one's norms. For the purpose of this paper, however, we can remain neutral on what are the conditions of adequacy. Thus, it may be that one's norms are epistemically adequate only if they are true, they are reliably truth or knowledge conducive, they lead to internally consistent beliefs, and so on. The important point for the current discussion is that, if one is rationally permitted to trust oneself in the way described, then, whatever conditions do govern the epistemic adequacy of epistemic norms, those conditions will be satisfied. The second condition is straightforward enough. The third condition is important since, even if the first two conditions are met, such reliability as to come from trusting oneself in the way described could still be a matter of accident. Yet, if one's perceptions conform to the relevant norms only by accident, one's self-trust would be mis-placed, not well-founded. That being so, there must be something that connects one's reliable conformity to the prescriptions of the relevant norms to the content of those norms. My suggestion is that this grounding will (typically) come via just the process of inculturation and internalisation of norms described.

This completes my description of the inculturation into trust model of norm-following. With this model in hand, we can return to consider the question of how we ought to respond to deep disagreement.

4. Responding to Deep Disagreement

In this section, I return to the normative problem of deep disagreement. In section 1, I suggested that deep disagreements can be characterised as disagreements that are explained by the disputants following different epistemic norms. In section 3, I suggested that it is apt to describe the way in which we (typically) come to follow and

¹⁷Self-trust of this kind is specific to epistemic practices that involve the following of epistemic norms. Epistemic self-trust *in general* is often conceived in terms of one's fundamental reliance upon one's cognitive faculties and epistemic practices (e.g. Foley 2001; Enoch 2010; Schafer 2015; Fricker 2016). It is worth noting that the account of norm-following given here is consistent with this way of conceiving epistemic self-trust. This is because norm-following is plausibly an aspect of a wide range of our epistemic practices. Thus, the account here can be seen simply as (i) providing a story of how we come to have and employ those practices that involve norm-following and (ii) adding detail to the idea of what it means to rely on – and so trust – oneself in respect to those practices. Thanks to an anonymous reviewer for pushing me to expand upon this connection.

implement our epistemic norms as the manifestation of a certain kind of socially grounded epistemic self-trust. (Since this is a mouthful, from here on I'll talk more simply of *trusting one's norms*) Taken together, then, the normative problem of deep disagreement amounts to whether and when it is permissible in the context of deep disagreement to maintain trust in one's norms.

As per discussion in section 1, many deep disagreements will be such that the disputants are not in a position to recognise that their disagreement is deep. Call these *opaque* deep disagreements. Opaque deep disagreements are a relatively simple affair. If disputants are not in a position to recognise that the disagreement is explained by a difference in norms and they would be permitted to trust their norms outside of the context of disagreement, it is not the case that they ought to doubt themselves in respect to those norms in that context. Rather than questioning themselves in respect to their norms, then, they should look for a more *ordinary* explanation of the disagreement – e.g., differences in evidence, or a performance error on one of their parts – and they should respond as is appropriate for what they find. In other words, questions about how disputants ought to respond to opaque deep disagreements will be settled by whatever is the correct account of ordinary disagreement and what that account prescribes in the relevant circumstances.¹⁸

Whilst many deep disagreements will be opaque, in at least some cases we would expect the disputants to recognise the nature of their disagreement. Call such disagreements *transparent* deep disagreements. In principle, there is no reason that either domain-specific or disagreements involving more general norms cannot be transparent – this includes cases such as CONSPIRACY and DOCTORS.¹⁹ That being so, we still need an explanation of how disputants ought to respond to transparent deep disagreement. More specifically, as per the normative problem of deep disagreement, we need an explanation which allows that, were each disagreement transparent, it would be permissible for Tessa in CONSPIRACY to remain steadfast, whilst Empany and Evitt in DOCTORS ought to be conciliatory.

What I want to suggest, then, is that this is a matter of whether it is *reasonable to expect* one to see one's rival's norms as a genuine epistemic alternative to one's own. This, in turn, is a matter of whether one has a sufficiently rich stock of shared social epistemic experiences with followers of one's rival's norms of the kind that typically undergird the internalisation of epistemic norms and coming to trust those norms. If one does, then, one ought to doubt one's norms. As such, one ought not to continue trusting oneself in respect to those norms. Thus, one ought to be conciliatory in the face of the disagreement. If one does not have such resources, then, it is not reasonable to expect one to see one's rival's norms as a genuine epistemic alternative to one's own and so the disagreement does not give one reason to doubt one's own norms – even if the disagreement is transparent. That being so, it will be permissible to maintain trust in one's norms and so steadfastness will be appropriate.²⁰

¹⁸Since opaque disagreements do not require disputants to reject norms, then, the appropriate response will be either of the steadfast or mixed responses. This cashes out the suggestion from section 2 that the mixed response may be appropriate in some cases where disputants are not in a position to realise that their disagreement is deep.

¹⁹An anonymous reviewer points out that, depending on the nature of the disputed norms, a deep disagreement might be transparent, whilst disputants' fundamental norms remain opaque.

²⁰There is some similarity here to Boghossian's (2006, 2008) proposal that we are entitled to our norms until encountering an alternative that raises legitimate doubts. In that light, this section could be read as

Let's cash this out. First, I'll discuss what it is to see one's norms as a genuine epistemic alternative to one's own. Second, I'll discuss the conditions under which it is reasonable to expect one to see one's rival norms as a genuine alternative to one's own. Third, I will put this discussion into contact with our two examples.

As per the account given earlier, following epistemic norms is ordinarily a matter of being inculcated into a normatively loaded way of perceiving the world and trusting one's perceptions of the world to reliably track the application conditions and prescriptions of those norms. What I want to suggest here is that to see one's rival's epistemic norms as a genuine epistemic alternative to one's own is a matter of being able to imagine perceiving the world in just this way whilst following that person's norms. We might think of this as one's being able to manifest a kind of virtual self-trust that is grounded in one's ability to imaginatively inhabit a way of perceiving the world different from one's own. The greater the difference in norms the greater the leap of imagination required. This kind of imaginative act *may* be aided by the disputants reflecting upon the evidence they have that a disagreement is deep – but such intellectual consideration will not suffice to make the leap wholesale. Rather, much as I have suggested we typically come to internalise our own norms through exposure to the practices of others, this kind of leap would typically require an agent to have a rich stock of shared social-epistemic experiences with followers of the rival norm on which to draw. These are the kinds of experience that involve the recognition of epistemic authority, relationships of dependence and trust, and epistemic cooperation and collaboration; i.e., just the kind of resources that will typically ground an agent's coming to follow and trust in their own epistemic norms.

Importantly, whilst this kind of imaginative act involves being able to reliably track the application conditions and prescriptions of one's rival's norms, it does not consist of that alone. What is crucial is that one be able to do so in the manner of trusting norms. This is important since there are ways of tracking one's rival's norms other than the kind of imaginative act described above. Most obviously, one could simply pay attention to what one's rival believes and in what conditions they believe as they do. Merely tracking the prescriptions of competing norms – this way or another – is not, however, enough to afford one reason to doubt one's own norms. This is because knowing what a norm prescribes is not reason to think that those prescriptions *should* be followed or that the corresponding norms are candidates for being *good* norms. In contrast, when one comes to *see* one's rival's norms as a genuine epistemic alternative to one's own in the way described, one grasps – not only what the prescriptions of those norms are – but what it would be like to be normatively compelled to comply with those prescriptions. Thus, mimicking the way one develops a genuinely normative and epistemic sensibility re one's own norms as described in section 3. And, it is this grasping of how the prescriptions of one's rivals norms could be normatively compelling – just as are the prescriptions of one's own – in respect to which it can be appropriate to doubt oneself in respect to one's own norms.²¹

This brings us, then, to the question of when is it reasonable to expect one to see one's rival's norms as a genuine epistemic alternative to one's own – i.e. as per above, when the realisation that a disagreement is deep will afford one reason to doubt one's own norms. My suggestion here is that it will be reasonable to expect

giving an account of just when such doubts should arise. Thanks to an anonymous reviewer for drawing my attention to this comparison.

²¹Thanks to an anonymous reviewer for pushing me on these points.

this just when one shares the kind of social-epistemic experiences and interactions with followers of one's rival's norms that would *ordinarily* facilitate a *competent* follower of one's own norms in manifesting the virtual self-trust in one's rival's norms described above. If one does have such resources, then, it is reasonable to expect one to see one's rival's norms as a genuine epistemic alternative to one's own – in which case, I suggest, realisation of the disagreement will afford one reason to doubt one's own norms. If one does not have such resources, then, it is not reasonable to expect one to see one's rival's norms as a genuine epistemic alternative to one's own – in which case, I suggest, realisation of the disagreement will not afford one reason to doubt one's own norms.

(It is worth noting that reasonable expectations so understood correspond to what the ordinary competent follower of one's norms is in a position to do. We could imagine agents who are, respectively, particularly gifted or particularly incompetent when it comes to the kind of imaginative act described. I presume, though, that it is neither reasonable to expect disputants to be particularly gifted, nor that incompetence gives one a free pass when it comes to epistemic oughts.)

With this account in hand we can return to the question of how it can be permissible for Tessa in CONSPIRACY to remain steadfast, yet Empany and Evitt in DOCTORS ought to be conciliatory.

Take CONSPIRACY. In this case, the social dimensions of deep disagreement have a kind of inverted significance in respect to the rational status of Tessa's continued trust in her norms. That is to say that, given that the difference in norms is fundamental and wide, we would not expect Tessa and Connie to share the social-epistemic experiences and interactions that would ordinarily facilitate a competent follower of Tessa's norm's coming to see Connie's norms as a genuine epistemic alternative to her own. And, if that is so, on the account given here, it will not be reasonable to expect of Tessa that she see Connie's norms in this way. Thus, the realisation that the disagreement is deep will fail to undermine the rational status of her continued self-trust. Thus, even if the disagreement is transparent, Tessa can rationally respond in steadfast fashion.

As noted in section 2, the kinds of deep disagreement which threaten global scepticism tend to be of a kind with CONSPIRACY. Thus, the thought that it can be rational to maintain trust in one's norms even in cases where the difference in norms is transparent allows us to stave off the sceptical worries that are associated with these kinds of case.

The flip side to this suggestion is that the possibility that those party to transparent deep disagreement will be permitted to stay steadfast in a narrower, domain-specific case of deep disagreement such as DOCTORS will be comparatively lower. After all, as we saw in section 2, the contents of the rival's norms in this case are such that we would not expect followers of each to doubt each other's competence as a matter of course. Correspondingly, there is no reason to think that they would not have the rich stock of shared social experiences that might spur the imaginative leap to understanding the other's epistemic sensibilities that is required to see the other's norms as representing a genuine alternative to their own. If Empany and Evitt do have such resources, though, it is reasonable to expect them to see each other's norms as a genuine alternative to their own. In which case, they ought to doubt their norms. As such, the pair ought to respond to the realisation of the nature of the disagreement in conciliatory fashion. As noted earlier, such a response will still have revisionary consequences – though not the comprehensive and wide-ranging ones associated with cases such as CONSPIRACY. Again, though, this seems to match up well with the intuitions that led

us to consider such cases to start with and to consider whether a non-uniform approach to deep disagreement may be the correct approach to deep disagreement.

5. Comments and Clarifications

Before wrapping up, I shall briefly discuss two areas worth additional clarification.²²

5.1. Conciliation and the absence of understanding

As described above, it is reasonable to expect someone to see another's norm as a genuine epistemic alternative to their own when they are in a position to gain a certain kind of imaginative understanding of that norm. This is a slight simplification. This is because there can be cases in which one is not in a position to understand one's interlocutor's norms in the way described yet should see them as a genuine epistemic alternative to one's own – or even as *superior* – nonetheless. The most obvious and important of such cases are those in which disputants follow different norms because one is an *expert* relative to the content of the disagreement and the other is not. In such cases, the expert is likely to follow norms that the non-expert cannot imaginatively understand. Yet, at the same time (it seems that) the non-expert should, typically, recognise the expert *and* their norms as epistemically superior. At first blush, this seems to be in tension with the account given so far. However, this is not the case. The important point, I suggest, is that such a situation will occur only when the non-expert should think of the expert and their norms as superior *according to norms that they follow* OR according to norms that they *should see* as a genuine epistemic alternative to their own. In such cases, though they could not follow the expert's norms themselves, it is reasonable to expect the non-expert to have an *indirect* understanding of those norms. For, not only could they follow the prescriptions of the expert's norms by paying attention to what the expert believes but, crucially, they will be in a position to grasp how the prescriptions of those norms could be normatively compelling. How? Just because their own norms – or genuine alternatives – tell them they ought to think of the expert and their norms as their epistemic superior. In light of this additional detail, then, we can add to the original account that someone party to transparent deep disagreement should respond to the disagreement in conciliatory fashion not only, as discussed, when:

- (i) It is reasonable to expect them to see their rival's norms as a genuine alternative to their own *because* a competent follower of their own norms should grasp what it is like to follow those norms.

But also, when:

- (ii) It is reasonable to expect them to see their rival's norms as superior to their own *because* a competent follower of their own norms – or genuine alternatives – should think of those norms as superior to their own.

So much for the first clarification. Let's move to the second.

²²Thanks to an anonymous reviewer for pushing me to clarify the two issues discussed below.

5.2. Self-trust and deep disagreement, self-trust and ordinary disagreement

The second point worth clarifying is how the current self-trust based account of deep disagreement relates to self-trust based accounts of ordinary peer disagreement. More specifically it is worth explaining why the current account is not subject to the same objections as are self-trust based defences of (uniform-)steadfastness in situations of ordinary peer disagreement.

There are a variety of self-trust based accounts of ordinary peer disagreement. To keep things brief I shall focus on the most established, which draws upon the reliance model of self-trust mentioned earlier to argue in favour of steadfastness in the face of deep disagreement. Variations of this argument are offered by Enoch (2010), Wedgwood (2010), and Schafer (2015).²³ Roughly, the argument goes as follows:

1. All epistemic deliberation begins from reliance upon one's own cognitive faculties and epistemic practices.
2. Therefore, on pain of scepticism, I must have a default *entitlement* to trust those faculties and practices as well as the deliverances of those.
3. However, if I am *entitled* to trust myself in this way, then, in case of disagreement over-p with someone I previously judged to be my peer, I will be entitled to trust my prior belief about p when determining whether that person is *in fact* my peer in respect to the issue under dispute.
4. But if I can do that, I can justifiably downgrade my interlocutor from the status of peer on the basis of the disagreement.
5. In which case, I can justifiably remain steadfast in the face of disagreement with someone I previously took to be my peer.

The principal problem with this argument is that it overlooks the fact that there are (at least) two prior beliefs relevant to how I ought to respond to the disagreement: my belief about p *and* my prior belief that my interlocutor is my epistemic peer. Even if it is correct to say that I am entitled to trust myself when determining how to respond to the disagreement, then, the question of *which* belief I should trust remains open. But, if that is the case, then so too does the question of how I should respond to the disagreement. Thus, the argument in favour of steadfastness fails to go through (at least as presented here).

If the most established self-trust account of ordinary disagreement is flawed in this way, then, why is the same not true of the account of deep disagreement given here? Simply because the account of deep disagreement given here is not a defence of uniform-steadfastness about deep disagreement. Rather, the account here is a non-uniform one – intended to explain how in some cases steadfastness can be appropriate and in others conciliation. Crucially, then, the suggestion is not that one has a default entitlement to one's own beliefs about p in case of disagreement. Instead, we might take it to offer an account of when – in cases of deep disagreement – one is permitted to continue trusting oneself and when one has a defeater for any such trust. Just what is lacking on the self-trust defence of steadfastness in ordinary peer disagreement.²⁴

²³See Pasnau (2015) and Peter (2019) for discussion relating to less standard conceptions of self-trust.

²⁴Thanks to an anonymous reviewer for their thoughtful suggestions on how to improve this paper, as well as the editors of the volume. For comments on earlier drafts and help developing the research from which this paper draws, thanks are owed to T. Ryan Byerly, Charlie Crerar, Paul Faulkner, Miranda Fricker, and Trystan Goetze, as well as to colleagues in the philosophy department at the University of

References

- Boghossian P.** (2006). *Fear of Knowledge: Against Relativism and Constructivism*. Oxford: Oxford University Press.
- Boghossian P.** (2008). 'Epistemic Rules.' *Journal of Philosophy* **105**(9), 472–500. <https://doi.org/10.5840/jphil2008105929>.
- Brotherton R. and French C.C.** (2015). 'Intention Seekers: Conspiracist Ideation and Biased Attributions of Intentionality.' *PLoS ONE* **10**(5), e0124125. <https://doi.org/10.1371/journal.pone.0124125>.
- Christensen D.** (2009). 'Disagreement as Evidence: The Epistemology of Controversy.' *Philosophy Compass* **4**(5), 756–67. <https://doi.org/10.1111/j.1747-9991.2009.00237.x>.
- Christensen D.** (2011). 'Disagreement, Question-Begging and Epistemic Self-Criticism.' *Philosopher's Imprint* **11**(6). Corpus ID: 155271077.
- Elga A.** (2007). 'Reflection and Disagreement.' *Noûs* **41**(3), 478–502. <https://doi.org/10.1111/j.1468-0068.2007.00656.x>.
- Enoch D.** (2010). 'Not Just a Truthometer: Taking Oneself Seriously (but not Too Seriously) in Cases of Peer Disagreement.' *Mind*, **119**(476), 953–97. <https://doi.org/10.1093/mind/fzq070>.
- Faulkner P.** (2011). *Knowledge on Trust*. Oxford: Oxford University Press.
- Faulkner P.** (2016). 'Agency and Disagreement.' In P. Reider (ed.), *Social Epistemology and Epistemic Agency: Decentralizing Epistemic Agency*, pp. 75–90. London: Rowman and Littlefield.
- Feldman R.** (2009). 'Evidentialism, Higher-Order Evidence, and Disagreement.' *Episteme* **6**(3), 294–312. <https://doi.org/10.3366/E1742360009000720>.
- Fogelin R.** (1985). 'The Logic of Deep Disagreements.' *Informal Logic* **7**(1), 1–8. <https://doi.org/10.22329/il.v7i1.2696>.
- Foley R.** (2001). *Intellectual Trust in Oneself and Others*. Cambridge: Cambridge University Press.
- Fricke M.** (2007). *Epistemic Injustice*. Oxford: Oxford University Press.
- Fricke E.** (2016). 'Doing (Better) What Comes Naturally: Zagzebski on Rationality and Epistemic Self-trust.' *Episteme* **13**(2), 151–66. <https://doi.org/10.1017/epi.2015.37>.
- Gifford F.** (2017). 'The Biomedical Model and the Biopsychosocial Model in Medicine.' In M. Solomon, J.R. Simon and H. Kincaid (eds), *The Routledge Companion to Philosophy of Medicine*, pp. 445–54. New York, NY: Routledge.
- Holton R.** (1994). 'Deciding to Trust, Coming to Believe.' *Australasian Journal of Philosophy* **72**(1), 63–76. <https://doi.org/10.1080/00048409412345881>.
- Jones K.** (1996). 'Trust as an Affective Attitude.' *Ethics* **107**(1), 4–25.
- Jones K.** (2012a). 'The Politics of Intellectual Self-trust.' *Social Epistemology* **26**(2), 237–51. <https://doi.org/10.1080/02691728.2011.652215>.
- Jones K.** (2012b). 'Response to Daukas.' *Social Epistemology Review and Reply Collective* **1**(10), 1–7.
- Kappel K.** (2012). 'The Problem of Deep Disagreement.' *Discipline Filosofiche* **22**(2), 7–25.
- Kappel K.** (2021). 'Higher Order Evidence and Deep Disagreement.' *Topoi* **40**, 1039–50.
- Kelly T.** (2010). 'Peer Disagreement and Higher-Order Evidence.' In R. Feldman and T.A. Warfield (eds), *Disagreement*, pp. 111–74. Oxford: Oxford University Press.
- Lackey J.** (2010). 'A Justificationist View of Disagreement's Epistemic Significance.' In A. Haddock, A. Millar and D. Pritchard (eds), *Social Epistemology*, pp. 298–325. Oxford: Oxford University Press.
- Littlejohn C. and Turri J.** (2014). 'Introduction.' In C. Littlejohn and J. Turri (eds), *Epistemic Norms: New Essays on Action, Belief, and Assertion*, pp. 1–10. Oxford: Oxford University Press.
- Lynch M.P.** (2010). 'Epistemic Circularity and Epistemic Incommensurability.' In A. Haddock, A. Millar and D. Pritchard (eds), *Social Epistemology*, pp. 262–77. Oxford: Oxford University Press.
- Lynch M.P. and Silva P.** (2016). 'Why Worry about Epistemic Circularity?' *Journal of Philosophical Research* **41**, 33–52. <https://doi.org/10.5840/jpr20165362>.
- MacIntyre A.** (1977). 'Epistemological Crises, Dramatic Narrative and the Philosophy of Science.' *The Monist* **60**(4), 453–72.

Tartu. This research was supported by the University of Tartu ASTRA Project PER ASPERA, financed by the European Regional Development Fund. This paper is partly based on the author's PhD thesis, completed 2020, Sheffield University.

- Matheson J.** (2009). 'Conciliatory Views of Disagreement and Higher-order Evidence.' *Episteme* 6(3), 269–79. <https://doi.org/10.3366/e1742360009000707>.
- Matheson J.** (2015). *The Epistemic Significance of Disagreement*. New York, NY: Palgrave Macmillan.
- Matheson J.** (2021). 'Deep Disagreements and Rational Resolution.' *Topoi* 40, 1025–37. <https://doi.org/10.1007/s11245-018-9576-y>.
- McDowell J.** (1978). 'Are Moral Requirements Hypothetical Imperatives?' *Proceedings of the Aristotelian Society, Supplementary Volumes* 52, 13–42.
- McDowell J.** (1979). 'Virtue and Reason.' *Monist* 62(3), 331–50. <https://doi.org/10.5840/monist197962319>.
- McDowell J.** (1998). *Mind, Value, and Reality*. Cambridge, MA: Harvard University Press.
- Pasnau R.** (2015). 'Disagreement and the Value of Self-trust.' *Philosophical Studies* 172(9), 2315–39.
- Peter F.** (2019). 'Epistemic Self-Trust and Doxastic Disagreements.' *Erkenntnis* 84, 1189–205. <https://doi.org/10.1007/s10670-018-0004-x>.
- Pritchard D.** (2021). 'Wittgensteinian Hinge Epistemology and Deep Disagreement.' *Topoi* 40, 1117–25. <https://doi.org/10.1007/s11245-018-9612-y>.
- Ranalli C.** (2020). 'Deep Disagreement and Hinge Epistemology.' *Synthese* 197, 4975–5007. <https://doi.org/10.1007/s11229-018-01956-2>.
- Ranalli C.** (2021). 'What is Deep Disagreement?' *Topoi* 40, 983–98. <https://doi.org/10.1007/s11245-018-9600-2>.
- Schafer K.** (2015). 'How Common is Peer Disagreement? On Self-Trust and Rational Symmetry.' *Philosophy and Phenomenological Research* 91(1), 25–46. <https://doi.org/10.1111/phpr.12169>.
- Solomon M.** (2015). *Making Medical Knowledge*. Oxford: Oxford University Press.
- Sripada C.S. and Stich S.** (2012). 'A Framework for the Psychology of Norms.' In P. Carruthers, S. Laurence and S. Stich (eds), *Collected Papers, Volume 2*, pp. 285–310. Oxford: Oxford University Press.
- Wedgwood R.** (2010). 'The Moral Evil Demons.' In R. Feldman and T.A. Warfield (eds), *Disagreement*, pp. 216–46. Oxford: Oxford University Press.

Simon Barker earned his PhD from Sheffield University in 2020 before working at the University of Tartu in Estonia from 2020–2022 as ASTRA guest lecturer in theoretical philosophy. Currently he is teaching at the University of Tartu and researching on the intersections between epistemic self-trust, mental health and epistemic injustice.