

**PRINCIPLES OF DISAGREEMENT,
THE PRACTICAL CASE FOR EPISTEMIC SELF-TRUST,
AND WHY THE TWO DON'T GET ALONG**

Simon Barker

The University of Sheffield

Abstract. This paper discusses the normative structure of principles that require belief-revision in the face of disagreement, the role of self-trust in our epistemic lives, and the tensions that arise between the two. Section 2 argues that revisionary principles of disagreement share a general normative structure such that they prohibit continued reliance upon the practices *via* which one came to hold the beliefs under dispute. Section 3 describes an *affective* mode of *epistemic self-trust* that can be characterised as one's having an attitude of optimism about the reliability of the practices *via* which one forms beliefs. The availability of *affective self-trust* is crucial to being able to rationally resist doubts about those practices. However, as argued in the final section of the paper, following revisionary principles of disagreement can render such trust *psychologically unavailable*. Consequently, attempts to follow such principles can have doxastic consequences that predictably exceed their normative implications.

Keywords: Epistemic self-trust, epistemology of disagreement, epistemic trust, epistemic peers, affective trust, epistemic practices

DOI: <https://doi.org/10.3176/tr.2020.3.07>

1. Introduction

This paper explores tensions that arise between principles that require one to revise one's beliefs in the face of disagreement and the role of epistemic self-trust in our everyday intellectual lives.

Broadly speaking, epistemologists working on disagreement are concerned with questions about how we *should* respond to disagreement. In particular, we might consider whether there are any substantive epistemic principles of the kind that require one to revise one's beliefs in certain general conditions or classes of disagreement. On

a meta-level, we might consider also whether such principles conform to any more general normative structure. In sections 2.1 and 2.2, I survey the general terrain of the epistemology of disagreement, and especially the literature on peer disagreement, to address just this question. I argue that whatever else their content, all such principles are committed to proscribing, in the relevant circumstances, against one's continued reliance on the practices *via* which one came to hold the beliefs under dispute.

In section 3, I turn to the topic of epistemic self-trust. Broadly speaking, epistemic self-trust can be understood as the combination of reliance on one's epistemic practices and some conscious positive attitude towards those practices. In 3.1, I develop what I call the 'practical case for self-trust', according to which practical limitations upon the possibility of establishing the reliability of one's epistemic practices push us to acknowledge the existence of a mode of epistemic self-trust *without* belief. In 3.2, I draw on the wider literature on trust to explore the idea that epistemic self-trust without belief is *affective self-trust*, characterised as one's having an attitude of optimism about one's reliance on one's own epistemic practices. I wrap up this discussion by considering what kinds of consideration will determine and influence the *psychological availability* of *affective self-trust*.

In section 4, I put the discussions of the epistemology of disagreement and epistemic self-trust into contact to highlight significant tensions between revisionary principles of the kind discussed in section 2 and the psychological availability of affective self-trust. These tensions, I suggest, present a serious problem for accounts of disagreement committed to such principles.

2. Principles of disagreement

2.1. Questions

Disagreements are a ubiquitous feature of our social lives to which we can and do respond in a variety of ways: Sometimes, our response to disagreement is *conciliatory* – we lower our confidence in the disputed beliefs, suspend judgement on the issue, or even accede to the views of our interlocutors. Other times, we respond in more *steadfast* fashion, sticking to our beliefs despite realising that others have come to believe differently. Given the different ways in which we can and do respond to disagreement, the question arises: how *should* we respond to disagreement?

Given the ubiquity of disagreement, never to revise one's beliefs in the face of disagreement would be tantamount to a thoroughgoing dogmatism. Presuming that one ought not to be dogmatic, then, the real question is not whether one ought *ever* to revise one's beliefs in the face of disagreement, but whether those cases that do generate such normative demands conform to any substantive epistemic principles of the kind that track general conditions and classes of disagreement.

It is this latter question with which I engage in the first section. Rather than try to identify specific principles, however, I want to come at the question from above and consider whether such principles conform to any more general normative structure. As I shall argue, *substantive revisionary principles* may differ in terms of the conditions

in which they apply, and the degree and method of revision prescribed. However, all such principles will be committed to proscribing, in the relevant conditions, against continued reliance on the epistemic practices *via* which the agent in question came to hold the beliefs under dispute.

As the reader will no doubt be aware, epistemologists have, for the most part, considered questions about the normative significance of disagreement in respect to the class of disagreements between epistemic peers. Likewise, it is in discussion of peer disagreements that epistemologists have most explicitly posited substantive revisionary principles of disagreement. For that reason, this discussion will also be a useful starting point to consider questions about the general structure of such principles.

2.2. From peers to principles

To illustrate the idea of peer disagreement, consider the following example:

CASE-1: Kazimira is a highly competent journalist working at a respected newspaper. She is currently investigating the possibility of electoral fraud in the recent general election. Throughout her investigation Kazimira has been in contact with her colleague Tom, who is also looking into the issue. Tom is also a highly competent journalist and Kazimira recognizes that. Throughout their inquiries, each has shared all their research with the other. At the next editorial meeting, Kazimira presents her research and her conclusion that there was electoral fraud. To her surprise, Tom – who clearly seems to be in full command of his cognitive faculties at the time – expresses his disagreement. In Tom’s opinion, the evidence that he and Kazimira have collected does not support the conclusion that there was electoral fraud. As it happens, the evidence Tom and Kazimira collected supports Kazimira’s conclusion, not Tom’s, and Kazimira competently assessed the bearing of that evidence on the possibility of electoral fraud, Tom did not.

Consider this case from Kazimira’s perspective. Given the stipulations about how each of the pair has performed in their assessment of the shared evidence, there is no question here that, prior to Tom’s exclamation of disagreement, Kazimira is rationally permitted to believe there was electoral fraud and Tom ought not to deny this. Despite that asymmetry, however, Kazimira enjoys no clear epistemic advantage of the kind that would allow her to settle the disagreement in her favour *independently of the substance of the disagreement itself*. Nor, for that matter, does Tom have any such advantage that would allow her to settle the disagreement in his favour. Moreover, this symmetry holds whether one’s view is that the normative features of the case are determined by the objective facts about what the case-evidence supports and Tom and Kazimira’s competence when it comes to assessing that evidence (as may be so on *externalist* accounts of disagreement)¹; the evidence that Kazimira has about

¹ See Lasonen-Aarnio (2014) for arguments in favour of externalism in general as an approach to problems of higher-order evidence, including those associated with peer disagreement. More specifically, Lackey’s (2010, 2013) justificationist view includes elements of reliabilism; whilst Hawthorne and Srinivasan (2013) offer a knowledge-first account of disagreement.

their comparative competence and access to the case-evidence (as per *evidentialist* accounts)²; or Kazimira's beliefs about the case-evidence and their comparative competence (as per accounts that lean on *subjective* conceptions of rationality).³ Taken together, these features of the case have the result that Kazimira cannot permissibly *both* take a position on the question of who is more likely to be mistaken about the way in which the election was conducted *and* bracket from consideration of that issue the fact that she and Tom disagree.

Whilst the literature offers a number of more precise characterizations of the peerhood relationship, I would suggest that, generally speaking, these can be understood as encapsulations of this feature of Case 1 cashed out in terms of the author's own preferred normative framework.⁴ For the purpose of the current discussion, then, let's say that Case 1 is a case of peer disagreement, and, give a catch-all definition of epistemic peerhood as follows:

Epistemic peers

If S believes *p*, S realises that R disagrees with S about *p* and, *independently of the substance of the disagreement*, S ought not to believe that she has the relative advantage over R vis-à-vis *p*, or that R has the relative advantage over S vis-à-vis *p*, then R is S's **epistemic peer**.

Accounts of the normative significance of disagreement under such conditions can be divided into three camps. *Conciliationists* argue that S should always revise her beliefs in the direction of her peers' (e.g. Christensen 2009, 2011, 2016, Elga 2007, Feldman 2007, 2009, Matheson 2009, 2015). Defenders of *steadfastness* argue that, under such conditions, S can be permitted to retain her original beliefs (e.g. Enoch 2010, Schafer 2015, Wedgwood 2010).⁵ Those who advocate *non-uniform* views maintain that the appropriate response to disagreement under such conditions can vary between cases (e.g. Faulkner 2016, Feldman 2009, Kelly 2010, 2013, Lackey 2010, 2013).

Discussing the provenance of this divide, David Christensen keenly observes that:

[i] All parties hold that the proper response to learning of another's disagreement depends on one's epistemic evaluation of that person.

[...]

² See for example, Matheson (2009, 2015), Feldman (2006, 2007, 2009).

³ See, for example, Elga (2007), Enoch (2010).

⁴ For example, Lackey admits to being "enough of an externalist about justification to require that the process or faculty responsible for the production of the belief in question be reliable or otherwise appropriately truth-conducive" (Lackey 2010: 321). In line with which, we might suppose, disputants will only be peers if they *are* equally familiar with the evidence and equally capable at processing that evidence. Which, roughly, is just how Lackey defines peerhood.

⁵ Enoch, Schafer, and Wedgwood all lean on the idea that an entitlement to 'self-trust' can underwrite steadfastness in the face of disagreement. However, the way each conceives 'self-trust' is closer to mere reliance than the normative conception of self-trust I discuss in later sections, for that reason their discussions do not feature in that discussion.

[ii] The camps differ, though, on this question: In evaluating the epistemic credentials of an opinion expressed by someone who disagrees with me about a particular issue, may I make use of my own reasoning about this very issue? (Christensen 2011: 1. Numbering additional).

I think Christensen's assessment of where the literature stands on [i] is quite correct. Since the motivation for this is obvious enough, let's label this idea *credentials* and move on to [ii]. Call the question Christensen raises here *permissibility*. I presume Christensen's reference to 'reasoning' in *permissibility* is shorthand for a much broader range of ways *via* which we might come to hold beliefs. After all, we do not arrive at all our beliefs *via* practices of reasoning. We also form beliefs *via* perception, intuition, testimony and so on. Nonetheless, I take it to be a truism that, for any belief one might come to hold, there will be at least one such practice *via* which one came to hold that belief (let's call such practices 'epistemic practices'). With that caveat, I think Christensen's observation here is right on the money. For, not only do the camps differ on *permissibility* – it is *the* central issue when it comes to discussing the normative-epistemic dimensions of disagreement. Not least, as we shall see, when it comes to examining the shared structure of any substantive revisionary principles of disagreement. To see why, let's return to CASE-1.

Note first that, just as for any other belief that she might hold, Kazimira must have come to believe there was electoral fraud *via* some or other of her epistemic practices. For convenience, let's refer collectively to these as *journalism*. We have already seen that, prior to discovering the disagreement, Kazimira is permitted to believe there was electoral fraud. Given that she came to believe electoral fraud *via journalism*, then, it must also be so that, prior to realising the disagreement, Kazimira was permitted to rely on *journalism* – where we can understand that to mean that she was permitted to *act or reason* in a way that presupposes that *journalism* reliably yields truths over falsehoods.

Why think this? Because, whatever the other details of the correct metaphysics of belief, when one believes a proposition, one takes an attitude to that proposition such that it is more likely true than false. Kazimira's belief that there was electoral fraud, thus, presupposes the reliability of *journalism* in the sense that, *ceteris paribus*, were she to doubt that *journalism* is reliable, and since it is amongst the outputs of *journalism*, Kazimira ought also to doubt that the proposition 'there was electoral fraud' is more likely true than false. Correspondingly, it follows that since Kazimira is permitted to believe there was electoral fraud, prior to realising the disagreement, she must also be permitted to act or reason in such a way as to presuppose that *journalism* is reliable.

When we understand epistemic reliance in this way, however, it follows too that, if Kazimira continues to be permitted to rely on *journalism* after discovering the disagreement, then, so too will she be permitted to presuppose that there was electoral fraud when it comes to evaluating the credentials of Tom. That being so, Kazimira might reason along the following lines:

1. Tom believes 'there was no electoral fraud'.

2. There was electoral fraud.
3. Tom was subject to some significant performance error when he investigated the issue of electoral fraud.

In other words, if Kazimira is permitted to continue relying on *journalism* after realising the disagreement, Kazimira will be permitted to reason under the presumption that the disagreement is evidence of Tom's epistemic shortcomings. Thus, even assuming *credentials*, and though Tom is her epistemic peer as defined, Kazimira will be under no normative pressure to revise her belief that there was electoral fraud.

In contrast, if Kazimira is not permitted to continue relying on *journalism* after realising the disagreement, then, she ought to revise her belief that there was electoral fraud. Why? Because, given the details of the case, Kazimira has no support for believing that she has the epistemic advantage over Tom vis-à-vis electoral fraud that does not come via *journalism*. Thus, Kazimira could not permissibly believe that she has the advantage over Tom without relying on *journalism* to do so. And so, if her response to the disagreement is not to violate *credentials*, and she is not permitted to rely on *journalism*, Kazimira ought to revise her belief that there was electoral fraud.

As the case goes, then, the proper response for Kazimira to take to her realisation of Tom's dissenting opinion is determined by whatever is the correct answer to *permissibility* in the case. Since nothing about this analysis depends upon the identity of the disputants or the content of the disagreement, we can extend these observations to arrive at the following general rules of disagreement:

Ceteris paribus, where K is some belief-forming practice of S's, S comes to believe *p* via K, and S realises that R disagrees about *p*:

- (Rule-1) If S is permitted to rely on K upon realising the disagreement, S will be permitted to stay steadfast in believing *p*.
- (Rule-2) If S ought to revise her belief that *p*, S is not permitted to rely on K upon realising the disagreement.
- (Rule-3) If R is S's epistemic peer and S is not permitted to rely on K upon realising the disagreement, S ought to revise her belief that *p*.
- (Rule-4) If R is S's epistemic peer, and S is permitted to stay steadfast in believing *p*, S is permitted to rely on K upon realising the disagreement.

In line with Christensen's comments, then, we can observe that any position one might take on the proper response to peer disagreement is fully constrained by the stance one takes to the question of *permissibility*. Notice, however, that (Rule-3) and (Rule-4) follow from (Rule-1) and (Rule-2) *only* because of the way in which peer disagreement has been defined. Thus, where disputants are not peers in the strict sense defined, settling the question of *permissibility* will not fully settle the question of how the disputants ought to respond to the disagreement. Nonetheless, the observant reader will note that the more general (Rule-1) and (Rule-2) are not

specific to peer disagreement. Rather, these two rules apply to *any* conditions under which disagreement might occur. This point is crucial to our wider discussion, so it will pay to illustrate it with another example. Consider, then, an extension of CASE-1, imaginatively labelled CASE-2:

CASE-2: Following the editorial meeting Kazimira shares all her evidence on the election with her colleague Kristen. Kristen is a significantly more experienced and competent journalist than Kazimira and Kazimira recognizes this. Having carefully examined Kazimira's evidence, Kristen tells Kazimira that she doesn't see any evidence of electoral fraud.

In CASE-2, we might say, Kazimira disagrees with her *epistemic superior*. Drawing on the definition of peerhood from earlier, we might define what it is for one disputant to be another's epistemic superior as follows:

Epistemic Superior

If S believes *p*, S realises that R disagrees with S, and, *independently of the substance of the disagreement*, S ought to believe that R has a significant epistemic advantage over her vis-à-vis *p*, R is S's **epistemic superior**.

Plausibly, whatever we think of CASE-1, Kazimira is under considerably greater pressure to revise her belief that there was electoral fraud when she finds herself in disagreement with her epistemic superior. Echoing what we said about CASE-1, however, if Kazimira retains permission to rely on *journalism* upon realising that Kristen disagrees with her, it does not matter that Kristen is her epistemic superior. For, similarly to CASE-1, Kazimira would be permitted to reason as follows:

1. Kristen believes 'there was no electoral fraud'.
2. There was electoral fraud.
3. Kristen was subject to some significant performance error when she assessed the evidence for electoral fraud.

Just as in CASE-1, then, if Kazimira retains permission to rely on *journalism*, Kazimira can reason under the presumption that the disagreement is significant evidence against Kristen – just as we saw she would be able to do so against Tom in CASE-1. And so, even assuming *credentials*, Kazimira would be under no normative pressure to revise her belief that there was electoral fraud.

In other words, (Rule-1) applies in CASE-2 just as it does in CASE-1. And, since it is just the contrapositive of (Rule-1), so does (Rule-2). Thus, it also holds that: if Kazimira ought to revise her belief that there was electoral fraud in the face of disagreement with Kristen, then, Kazimira cannot be permitted to continue to rely on *journalism* in those circumstances.

What I would suggest CASE-2 illustrates is that (Rule-1) and (Rule-2) generalize to all cases and classes of disagreement. Correspondingly, then, and returning to our main theme, (Rule-1) and (Rule-2) also generalize to *any and all* substantive revisionary principles of disagreement. So, we might say of such principles:

Where *C* describes the context in which some general class of disagreement is instantiated, and *S* has come to believe *p* via *K*:

- (i) For any **substantive revisionary principle (SRP)** *T* that entails that *S* ought in some way to revise her belief *p* in *C* (in the direction of believing $\sim p$),
- (ii) *T* entails that *S* is not permitted to rely on *K* in *C*.

Notice that *SRPs*, so described, prescribe revision only within a certain context, i.e. *C*. This is not to say, however, that there cannot be plausible diachronic or practical *strategies* for responding to disagreement that can result in one's coming to believe conflictingly to one's original beliefs – and that involve relying upon *K* to do so. This would be so, for instance, on a two-stage strategy requiring *S* to suspend judgement on *p* within *C* (stage-1) and re-deploy *K* to reassess the issue when she has exited *C* (stage-2).⁶ At first blush, the plausibility of this kind of strategy appears to call into question the definition of *SRPs* given above. Crucially, however, a multi-stage strategy of this kind will only be plausible, if there are significant and relevant epistemic differences between the context in which *S* is required to suspend judgement on the original output of *K* (i.e. *p*) and the context in which she redeploys *K* to reassess the issue. In the absence of such a difference, on the other hand, there will be no explanation of how stage-1 could require *S* to suspend judgement on *p*, yet, as per stage-2, *S* be permitted to employ *K* to reassess the issue.

On any plausible variant of the multi-stage strategy, then, it is only *S*'s initial response at stage-1 that will be governed by the relevant *SRP*. In contrast, whatever response is required of *S* as a result of her redeployment of *K* will not be governed by that principle – just because that principle proscribes against *S* relying upon *K* until she has exited *C*. The plausibility of multi-stage strategies of this kind does not, then, call into question the entailment between *SRPs* and reliance after all – even if such strategies may result in one's coming to believe conflictingly with one's original, disputed beliefs.⁷

With that issue addressed, we can note that *SRPs* will still differ in the conditions under which they apply, and the degree and method of revision required. (For instance, one *SRP* may say that *S* ought to suspend judgement on *p* in *C*, a rival *SRP* may say that *S* ought to concede to her interlocutors and come to believe $\sim p$ in *C*). Nonetheless, as discussed, any *SRP* that applies to *C*, must proscribe against continued reliance on *K* in *C*. For, if not, then, when finding herself in *C*, *S* will be permitted to reason as if the disagreement is evidence against her interlocutors, and on that basis stay steadfast in believing *p*. It is this feature of *SRPs* that I will argue later in the paper generates tensions with epistemic self-trust and its role and importance in our epistemic lives. Let us turn now, then, to the topic of epistemic self-trust.

⁶ For instance, Levy (2020: 10) suggests that it would be appropriate for members of a community of experts to employ this kind of strategy if they encounter a dissenter from a widely held consensus view within that community.

⁷ Thanks to an anonymous reviewer for pushing me to consider such multi-stage approaches to disagreement.

3. Epistemic self-trust

3.1. Why self-trust?

Discussions of epistemic self-trust often begin from the observation that, as believers, we are each unavoidably dependent upon our own epistemic practices. So, for example, Tanesini writes that:

Some kind of pre-reflective self-trust must already be in place if we are to carry out any inquiry. Without it we would be paralysed. [...] Epistemic self-trust, therefore, consists of our propensities to *rely* on those aspects of our cognitive lives [...] that make our epistemic inquiries possible (Tanesini 2020: 222. Italics added).

And Fricker that:

Each one of us in one's everyday life *relies* on one's core package of cognitive faculties – perception, proprioception, memory, intellectual intuition and introspection – reliably to deliver one true beliefs. Once one appreciates that this reliance is, perforce, epistemically ungrounded [...], it is apt to call this ungrounded reliance epistemic self-trust (Fricker 2016: 154. Italics added).

We have already discussed aspects of the relationship of reliance between one's beliefs and one's own epistemic practices. What I want to consider in this section is why we might move from recognition of this relationship to the claim that epistemic self-trust has a fundamental role in our epistemic lives. Having done so, I shall move on to consider the question of how we should conceive epistemic self-trust if it is to fulfil that role. Echoing Jones's (2012a, 2012b) work on the topic of self-trust and drawing on literature on interpersonal trust more broadly, I will suggest that the fundamental form of epistemic self-trust is a genuine form of trust that can be characterised as one's having an affective attitude of *optimism* towards one's reliance on one's own epistemic practices.

Let us start by returning to the topic of reliance. Fortunately, we have already covered the key points and definitions. We can restate these as follows:

- Where K is amongst S's epistemic practices, S **relies** epistemically on K *iff* S acts or reasons in such a way that S's doing so *presupposes* that K reliably yields truths over falsehoods.
- For any belief *p* held by S, there will be some epistemic practice K upon which S relied in coming to believe *p*.
- *Ceteris paribus*, if S comes to believe *p* via K, and S is permitted to believe *p*, S is permitted to rely on K.

Whilst we could say more about the nature of epistemic practices, these three features of our reliance on those practices provide all we need to make the case for self-trust. As I see it, there are two ways that we might make that case: the

philosophical case, founded on consideration of contexts in which one is engaged in the distinctly epistemological project of establishing the reliability of one's epistemic practices; and the *practical case*, founded on consideration of the more ordinary contexts that we encounter in our everyday intellectual lives. Given that the focus of this paper is upon tensions between principles of disagreement and the role of self-trust in our everyday intellectual lives, I will focus on the practical case.⁸

3.1.1. *The practical case for epistemic self-trust*

Debate around the normative significance of disagreement rests on the presumption that we need not be global sceptics; moreover, it typically proceeds on the assumption that sweeping, if not global, sceptical consequences are also to be rejected. In the current context, then, we can assume that, for most of us, a significant number of the beliefs that we hold, across a wide range of topics, we are permitted to hold. Given what we have said before, this implies that most of us must be permitted to rely on at least some of our epistemic practices a significant amount of the time and across a wide range of topics. Presuming, then, that K is amongst those practices that S is particularly dependent upon in the way described, we can say that:

- (1) *Ceteris paribus*, S is permitted to rely on K a significant amount of the time and across a wide range of topics.

What I want to suggest first is that S could be permitted to rely on K in this way *even if* she has not established that K is reliable, and indeed, even if she has never considered whether K is reliable. How so? Simply because the project of establishing the reliability of the epistemic practices *via* which one formed a specific belief or set of beliefs is a distinctly epistemological project of the sort that we do not and cannot conduct in ordinary contexts. It would require, at the least, (i) identifying what the relevant practices are, (ii) identifying what kinds of evidence are relevant to establishing the reliability of those practices, (iii) collecting that evidence, and (iv) assessing that evidence – not to mention navigating the various theoretical problems that this kind of project might throw up. Yet, most of us are not trained to conduct such projects, most of us do not have the theoretical or conceptual resources to conduct such projects, and, perhaps most significantly, most of us do not have the time and practical resources to conduct such projects – at least not under normal circumstances. On pain of scepticism, then, rationality cannot in ordinary circumstances require the average person to have conducted such a project – or, indeed, even considered conducting such a project. That being so, we can move from (1) to:

- (2) *Ceteris paribus*, S is permitted to rely on K a significant amount of the time and across a wide range of topics, without having available the resources that would allow S to establish that K is reliable.

⁸ Roughly, the *philosophical case* proceeds by noting that attempts to fully establish the reliability of one's basic epistemic practices will be subject to epistemic circularity. On pain of scepticism, then, or so the argument goes, we must be permitted to *trust* those practices without justifying belief to that effect. See Foley (2001), Fricker (2016), and Zagzebski (2012) for arguments along these lines. See Alston (1986) for the seminal discussion of epistemic circularity.

Following Burge (2013) we can refer to this relationship in terms of entitlement.

[E]ntitlements are epistemic rights or warrants that need not be understood by or even accessible to the subject. We are entitled to rely, other things equal, on perception, memory, deductive and inductive reasoning, and on – I will claim – the word of others. [...] Philosophers may articulate these entitlements. But being entitled does not require being able to justify reliance on these resources, or even to conceive such a justification (Burge 2013: 230).

Thus, we might restate (2) as:

- (3) *Ceteris paribus*, S is entitled to rely on K.

A number of writers appear to accept this entitlement to rely on our epistemic practices as itself constitutive of an entitlement to self-trust (e.g. Foley 2001, Fricker 2016).⁹ I would suggest, though, that this is only the first staging post in the case for self-trust. ‘Trust’, as it is typically understood in the literature, includes some conscious attitudinal component taken by the trusting toward the trusted. Thus, an entitlement to rely is an entitlement to self-trust in name only.¹⁰ To make the case for a genuine form of self-trust, then, we need also to consider contexts in which that entitlement does not hold. More specifically, and to return to S, we need to consider contexts in which both:

- (4) S realises, or ought to realise, that she came to believe *p* via K,
 (5) S realises, or ought to realise, that beliefs formed *via* K are no more likely to be true than false if K is not reliable.

Let’s call contexts like (4) *contexts of realisation* and contexts like the conjunction of (4) and (5) *contexts of doubt*. Given (4) and (5), S is confronted directly with the question ‘Is this practice K, *via* which I formed the belief *p*, reliable?’ And, if S cannot, or ought not to, give a positive answer to that question, S ought to rely on K no longer. Why? Because, to be faced with the question of whether K is reliable, yet take no attitude toward K that would involve the presumption that K is, or is likely, to be reliable, *just is* to doubt that K is reliable. And, if S realises that she came to believe *p via* K and S doubts that K is reliable, *ceteris paribus*, S ought not to continue believing *p*. Thus, given (4) and (5) it follows that:

- (6) S ought to take a positive attitude toward her reliance on K, or not rely on K.

⁹ As noted earlier, Enoch, Schafer, and Wedgwood all treat self-trust in this way in their respective discussions of peer disagreement.

¹⁰ Fricker explicitly makes the claim that it is apt to refer to a relationship of ‘ungrounded reliance’ as a relationship of *trust*. It seems to me, though, that, even if we accept this claim, it will only motivate a further distinction, between, say, *attitudinal* and *non-attitudinal* trust. To my mind, the thesis of self-trust is most interesting when it is a thesis about some form of *attitudinal trust*.

The question, then, is what kind of attitude would underwrite S's continued reliance on K, when faced with the question of K's reliability?

The obvious answer is that S might close this question *via* belief – specifically belief to the effect that 'K is reliable'. The combination of reliance and belief in the reliability of the thing relied upon describes a genuine form of trust. Following Faulkner (2011) we can refer to this as *predictive trust*. Correspondingly, were S to come to believe 'K is reliable' and so continue to rely on K, S could be said to have *predictive self-trust* in respect to K. We can define *predictive self-trust* as follows:

Predictive self-trust (PST)

Where K is amongst S's epistemic practices, S has **predictive self-trust** (in respect to K) *iff*:

- (i) S relies on K
- (ii) S believes that 'K is reliable'.

Given the belief component, we can add also that S will be permitted to have *PST*, *iff* S is permitted to believe that 'K is reliable'. Thus, it follows that, if S is permitted to believe that 'K is reliable' in a *context of doubt*, K will also be permitted to rely on K in that context. Even accepting this, however, we do not arrive at the conclusion that epistemic self-trust is of special importance in our everyday intellectual lives. After all, we have already seen that establishing the reliability of one's epistemic practices requires undertaking a sophisticated and distinctly epistemological research project. And, as we also saw, there are significant practical limitations upon our capacity to carry out such projects in anything like ordinary circumstances. *PST*, then, will have no special role to play in our everyday intellectual lives, just because there are significant practical limitations upon the availability of *PST*.

Once we have in mind the idea that S's situation requires a form of epistemic self-trust, however, the question of whether there are forms of self-trust other than *PST* comes into play. For, just as predictive trust is only one among various forms of trust more generally – so it may be that *PST* is only one among various forms of epistemic self-trust. And it is in this light that the points about the limitations of *PST* are not without significance when it comes to our more general discussion. For notice that, if *PST* is the only form of epistemic self-trust, then, by way of the same practical considerations that limit the availability of *PST*, it will be so that either (i) we rarely encounter contexts of doubt in our everyday intellectual lives, or (ii) we should accept a thoroughgoing and wide-scope, if not global, scepticism. On the presumption that no such scepticism is required in the world we live in, then, the practical case for epistemic self-trust rests upon the question of whether contexts of doubt are regular features of our everyday intellectual lives. I would suggest they are. To illustrate this thought, consider the following examples:

- a. Juri is a high-school teacher and is explaining to a colleague that he gave one of his students an extension on their homework because they told him that their internet had been out for the week. The other teacher just laughs and says 'Juri, you are just too gullible!'

- b. Mischa is a hospital resident on the tail end of a particularly long shift. After deciding that her current patient needs more morphine, she finds herself yawning and rubbing her eyes.
- c. Samantha is a keen user of social-media sites on which she follows a wide range of social-justice focused accounts and profiles. Watching a news report on #BlackLivesMatter protests in the USA, Samantha, who is white and upper-middle class, thinks to herself ‘at least the UK has no real problem with racism’.

I take it that none of these examples are extraordinary. Yet surely, Juri, Mischa, and Samantha should all have some sense that the question of whether the relevant practices are reliable has become salient. After all, Juri’s colleague directly questions Juri’s judgements of who to trust; Mischa’s yawning and rubbing her eyes are potentially symptoms of exhaustion; and someone in Samantha’s social position should *without doubt* question whether her intuitions about race relations in the UK are on point. In other words, despite none describing extra-ordinary circumstances – in all three examples, the individual in question encounters a context of doubt.

What I want to suggest, then, is that these examples attest to the ease with which contexts of doubt can, and most likely do, regularly arise in each of our everyday intellectual lives. But, if I am right about that, then, we should acknowledge the significant role that epistemic self-trust has to play in our everyday intellectual lives. For, if contexts of doubt are a regular feature of our epistemic lives, then, on pain of scepticism, we should allow that there is some form of epistemic self-trust that is not subject to the same practical limitations as *PST* (whether or not it could plausibly be available for Juri, Mischa, or Samantha). In other words, we should allow that there is a form of epistemic self-trust *without* belief. In the next section, I outline the general features of what I think such self-trust might be.

3.2. *Affective self-trust*

In the previous section, I presented the ‘practical case for self-trust’. As it turned out, we saw that the practical case is in the end a case for self-trust *without* belief. In this section, I want to introduce what I think is the best way of conceiving such self-trust. For reasons of space, this discussion will be more exploratory than critical, hopefully the reader will bear with me, nonetheless.

To get to the conception of self-trust I am interested in, it will be useful first to consider how the attitude of trust has been discussed in the broader corpus of work on this topic, particularly in discussion of interpersonal trust. Particularly promising here is the suggestion that genuine interpersonal trust has an *affective* dimension. One way that we might draw out this affective dimension is to place emphasis upon the *affective response* that the *trusting* is likely to experience when the *trusted* fails to make good on that trust. Along these lines, Richard Holton writes:

In cases where we trust and are let down, we do not just feel disappointed, as we would if a machine let us down. We feel betrayed. [...] [B]etrayal

is one of those attitudes that Strawson calls the reactive attitudes. These are attitudes that we normally take only towards people. We feel hurt or resentful when they let us down; grateful, perhaps touched, when they help (Holton 1994: 66).

Similarly, Faulkner, distinguishing ‘affective trust’ from ‘predictive trust’, writes:

The contrast [with predictive trust] is between expecting that something will happen and expecting something of someone. When we expect something of someone we are susceptible to certain reactive attitudes if they do not do what is expected [...] The general reactive attitude in play here is resentment: in *affectively trusting* S to ϕ , A will be prone to resentment were S to show no motivation to ϕ . This reactive attitude is the hallmark of the defeat of normative expectations (Faulkner 2011: 147).

To my mind, this emphasis upon the reactive attitudes associated with breakdowns in the trust relationship, captures something essential about interpersonal trust. Whilst, however, we may feel disappointed when our own epistemic practices let us down, arguably, it is neither apt nor coherent to describe our responses to these situations in terms of ‘betrayal’ or ‘resentment’. Thus, this dimension of interpersonal trust does not naturally extend to self-trust.

A more promising line of fit is offered by Karen Jones, who writes:

Trusting is composed of two elements, one cognitive and one affective or emotional [...] Roughly, to trust someone is to have an attitude of optimism about her goodwill and to have the confident expectation that, when the need arises, the one trusted will be directly and favourably moved by the thought that you are counting on her (Jones 1996: 5).

Here, instead of the nature of our responses to how trust relationships play out, Jones emphasises the affective nature of the attitude held *when* trusting. And crucially, whilst, on Jones’s account, the attitude of optimism takes the other person and their goodwill as its object in the context of interpersonal trust, it need not in other contexts. Instead, just as one may *believe* that X will ‘reliably yield truths and not falsehoods’ when X is another person or when X refers to one’s own epistemic practices; so one may be *optimistic* about relying on X to ‘reliably yield truths and not falsehoods’ when X is another person *or* one’s own epistemic practices. It is in this light, then, that I want to suggest, just as there is *affective interpersonal trust*, there is also an affective mode of (epistemic) self-trust. Borrowing from Jones, we can define this as follows:

Affective Self-Trust (AST)

Where K is amongst S’s epistemic practices, S has **affective self-trust** *iff*:

- (i) S epistemically relies on K

- (ii) S is *optimistic* about relying on K.¹¹

The key question, then, is what it means for S to be optimistic about relying on K.

Optimism, as I conceive it, is a primitive attitude that manifests in an agent's willingness to act in certain ways and willingness to think well of the person, practice, or relationship to which that action corresponds. In the context of trust and reliance, this attitude manifests in the willingness to rely on X to F and a willingness to think well of X and one's reliance on X. To think well of X in the sense implied involves making a series of presumptions in the context of action that correspond to having a positive affective and cognitive attitude toward X's F-ing. What I would like to suggest, then, is that for S to be optimistic about her reliance on K is for S to manifest a willingness to rely on that practice under the presumptions that:

- (i) It will be beneficial to her epistemic aims and interests if K reliably yields truths over falsehoods
- (ii) S will accrue those benefits if she acts as if K is reliable, and so
- (iii) K will reliably yield truths over falsehoods.

These presumptions allow *AST* to play a similar role in the rational structure of the agent's attitudes as belief does in the case of *PST*. Crucially though, they do not involve or imply belief – explicit or otherwise. Instead, we can think of optimism here as an affectively loaded way of *perceiving* or *experiencing* the world structured by a commitment to act in accordance with the three presumptions described. As Faulkner, describing the rational structure of interpersonal trust, puts it:

[I]f any of these propositions were not accepted, it would cease to make sense to say that A trusted because accepting this set of propositions is an expression of A's attitude of trust and so a commitment of A's decision to trust. This is not to suggest that trusting someone to do something involves explicitly committing to these propositions in one's reasoning. The claim is rather that the acceptance of these propositions partly defines how it is that the attitude of affective trust involves *seeing things in a certain light* (Faulkner 2011: 152. Italics added).

One way that we might describe this way of experiencing or perceiving one's reliance upon one's own epistemic practices, is such that, when S manifests self-trust in K, S has a *first-person* experience of that relationship as one in which she is already invested *and* that is worth investing in despite how doing so leaves her vulnerable. This is in contrast to the *third-person* perspective that S would take toward her reliance on K were she to arrive at *PST* via the kind of epistemological research project that we saw would be needed to establish that 'K is reliable'.

¹¹ Jones herself extends the affective conception of trust to self-trust in her work on the latter's relationship to epistemic injustice. Thus, she describes self-trust as a domain-relative "attitude of optimism about one's cognitive competence" (Jones 2012a: 243). It goes without saying that my account is heavily indebted to Jones's work. However, my discussion of self-trust's rational structure owes much also to Faulkner's account of affective interpersonal trust.

So conceived, *AST* comes with its own rational structure baked in. That is to say that the presumptions that are constitutive of *AST* themselves answer the question: ‘Why trust yourself?’. And so, we might say, *AST* is self-rationalizing trust. That being so, it is in the nature of *AST* that it can be arrived at, and appropriately so, without having conducted the kind of epistemological project we saw to be required for *PST*. And, crucially, since *AST* entails reliance, it follows that when one is permitted to trust in this manner, and when one manifests such trust, one will also be permitted to rely on the practices in question. In this light, then, I would suggest that *AST* can play the role in our everyday intellectual lives that I argued must be filled in at least some contexts of doubt if we are to avoid a thoroughgoing and wide-scope scepticism.

Having discussed the rational structure of *AST* and how *AST* can play the role required in the normative structure of our beliefs, I want to wrap up this section by focusing on the question of what determines the *psychological availability* of *AST*.¹²

As I see it, the psychological availability of *AST* depends upon two broad factors: one’s psychological character, and one’s personal history of circumstance and experience.

The availability of *AST* depends on one’s psychological character because one may not be psychologically inclined to see one’s own epistemic interests as worth investing in in the way that is characteristic of self-trust’s *first-person* perspective. Similarly, one may not be inclined to see the situations one finds oneself in as situations in which there are benefits to be accrued by continued reliance on one’s epistemic practices, as is characteristic of seeing the world through the affective lens of *optimism* (in the context of self-trust). Correspondingly, when an individual is not inclined to see themselves, or the situations in which they find themselves, in this positive affective light – *AST* will not be available to that individual. The availability of *AST* depends on one’s personal history of circumstance and experience, then, just because circumstance and experience exert profound and enduring influence over one’s psychological state. And these changes can, in turn, make significant differences to the psychological availability of *AST*. Most importantly in the current context, that influence can easily exceed (or fall behind) what the third-person perspective about those circumstances would merit – not because of any rational failing on the part of the individual in question, but simply because it is within the nature of *AST* to be so affected.

It is this last point that we shall see is central to the tension between *AST* and accounts of disagreement committed to substantive revisionary principles. In that light, let us move on to explore those tensions in more detail.

¹² Were the aim of this paper to employ the concept of *AST* to answer in-full the question of how we *should* respond to disagreement, we would also need to consider questions about the *rational availability* of *AST*. Since, that is not the aim of this paper, however, we can put these questions aside. (Though, for critical discussion of a similar approach to disagreement, see Peter 2019). Thanks to an anonymous reviewer for highlighting the need to clarify this point.

4. Principles of disagreement and the psychological availability of epistemic self-trust

In the first two sections of this paper, I explored, first, the broad terrain of debate in the epistemology of disagreement, with special attention to the question of whether any substantive revisionary principles (*SRPs*) that that debate may throw up conform to a more general normative structure; and second, the role of epistemic self-trust in our everyday intellectual lives. We are now in a position to see how these topics interact and, as I shall argue, how they are in tension. The first suggestion I make will be relatively trivial. It leads, though, to a second and deeper concern that arises when we consider what influence one's following *SRPs* may have upon the psychological availability of, specifically, *affective self-trust* (*AST*). In short, I will suggest that we might reasonably expect the trust-related consequences of following any *SRP* to exceed the normative implications of that principle. This poses a significant problem for any account of disagreement committed to such principles.

First, the trivial point. Recall from the schema introduced in section 2, that, if *S* comes to believe *p* via *K*, any substantive revisionary principle *T* that requires *S* to revise her belief *p* in circumstances *C* will entail that *S* is not permitted to rely on *K* in *C*. Correspondingly, any account of disagreement that endorses *T* will entail the same proscription against reliance. As we saw in section 3, however, on pain of scepticism, we should recognise that, in general, if *K* is amongst *S*'s epistemic practices, then, *ceteris paribus*, *S* will be *entitled* to rely on *K*. We saw also that that entitlement will not hold if *S* is in a *context of doubt* vis-à-vis *K*. In such a context, *S* will be faced directly with the question 'Is *K* reliable?', and given that question, *S* will be permitted to rely on *K* iff *S* has (and is permitted to have) the kind of positive attitude toward *K* that might rationally close that question. In that respect, we saw the primary role of epistemic self-trust, and specifically *AST*, is to underwrite one's continued reliance on one's epistemic practices in at least some of the contexts of doubt encountered in our everyday intellectual lives.

Taking this all together, then, we can recognise two points of contact between the epistemology of disagreement and epistemic self-trust:

1. In so far as *T* entails that *S* is not permitted to rely on *K* in *C*, *T* entails that *S* is not *entitled* to rely on *K* in *C*. Thus, if *S* is in *C* and *T* is true, *S* is in a *context of doubt* in respect to *K*.
2. If *S* is in a *context of doubt* in respect to *K* and *T* entails that *S* is not permitted to rely on *K* in *C*, then, *T* entails that *S* is not permitted to have *AST* in respect to *K* in *C*.

In other words, not only will *SRPs* imply that one cannot, in the relevant circumstances, rely on the epistemic practices *via* which one arrived at the beliefs under dispute, such principles will proscribe against a *trusting response* to disagreement in those circumstances.

On its own, this is not necessarily a problem for accounts of disagreement that endorse such principles. Whilst I have presumed throughout this discussion that we

should have serious qualms about endorsing any view that has significant sceptical consequences, the rational connection between *SRPs* and *AST* does not necessarily lead us any further down that road than do the relevant *SRPs* on their own. The crucial point here is that the connection described above implies a strictly one-to-one relationship between cases in which *SRPs* proscribe revision and cases in which such principles proscribe against a trusting response to disagreement. Thus, even if some *SRP* were to apply to a significant class of everyday disagreement (which it seems to me will be the case for any interesting *SRPs*), the *rational* connection between *SRPs* and *AST* implies no greater sceptical consequences than those that will already follow from that principle. Whilst some accounts of disagreement committed to such principles may raise sceptical concerns, then, this aspect of the relationship between *SRPs* and *AST* adds no weight to those concerns. However, the rational connection between prescriptions to revise and proscriptions against the trusting response to disagreement does not capture the full extent of those interactions. To get the full picture, we also need to consider the relationship between the psychological nature of *AST* and the potential psychological effects of following any *SRPs* that are in play. It is here that, I would suggest, more significant worries arise. To illustrate, consider one last variant on Kazimira's predicament:

CASE-3

As it happens, the disagreement Kazimira experienced in CASE-1 is not an isolated incident. Kazimira frequently encounters disagreement in meetings of the editorial board – no matter the quality of her work. These include cases in which, independently of the substance of the disagreement, Kazimira does have the relative advantage over her interlocutor. But far more instances are not like this – and not only is it often so that Kazimira disagrees with one of her peers, but it is common for other members of the group to express their joint disagreement with her views, as it is also common for her presumed superiors to pitch in too. In short, when in the editorial meeting room, Kazimira is exposed to a systematic and overbearing culture of disagreement. Nonetheless, Kazimira is (still) a highly competent journalist, and, just as it was so for her views on electoral fraud, the opinions that she presents to the group are, more often than not, meticulously researched, competently reached, and reliably true.

Given the structure of this case, we might presume that any interesting principled account of disagreement (where the relevant principles are *SRPs*) will imply that Kazimira will frequently have to revise her beliefs in respect to the opinions and ideas she brings to the group – even though these views are reliably based upon her excellent journalistic work outside of the meeting room. Though such consequences would be unfortunate for Kazimira, we have seen already that, in itself, this presents no special problem for any account that would make such demands upon Kazimira. The problem arises, though, when we consider what it would mean for Kazimira if she were to follow the prescriptions and proscriptions of such an account. Given the previous points, we already know that doing so would entail that Kazimira will

rarely have a *trusting response* to the culture of disagreement that she experiences in the meeting room. Supposing that she does respond in this way, then, the relevant question shifts from the question of whether *AST* is rationally available to her – but what responding to the normative demands of the relevant principles in this way will do for the psychological availability of *AST* to Kazimira.

Recall that the psychological availability of *AST* is determined by (i) the psychological character of the agent in question and (ii) the agent's personal history of circumstance and experience. What I would suggest in this case, then, is that Kazimira's exposure to the overbearing culture of disagreement in the meeting room, combined with her repeated failure to see herself and her reliance on her epistemic practices in the positive light of *AST*, could, entirely predictably, have a significant, deleterious, and wide-ranging impact upon Kazimira's psychology – and, thus, to the psychological availability of *AST*. Indeed, we might expect this to be so whatever Kazimira's initial perception of herself and her own epistemic competence, and certainly no less so because these disagreements are with those she otherwise presumes to be her peers and superiors. Indeed, ultimately, we might expect that by following the relevant principles of disagreement in the editorial meeting room, Kazimira may well come to a point where it is not self-trust that is readily available to her in contexts of doubt, but self-*distrust*. Where, as Karen Jones describes it, self-*distrust* involves an agent's "being disposed to feelings of lack of self-confidence, hesitancy to assert, discounting one's own judgment especially in the light of conflict, and rumination on one's competence" (Jones 2012b: 5).

If, however, Kazimira were to reach the point of self-*distrust*, and perhaps before that, it seems to me that we might also expect that the doxastic consequences of Kazimira's following the relevant *SRPs* in the meeting room will far outstrip the normative implications of those principles. For example, it would not be surprising were Kazimira to begin to conciliate in disagreements with those that she should otherwise see as being in an epistemically inferior position to her; likewise, it would not be surprising were Kazimira to begin to lose confidence in her beliefs on topics and issues in domains other than those discussed in the meeting room. Indeed, if the experiences in the editorial room are bruising enough, and the psychological trauma goes deep enough, it would not be surprising to see Kazimira revising her beliefs across wide swathes of her intellectual life, even in areas entirely unconnected to her experiences in the meeting room. And crucially, all of this may not be surprising, and some of it may even be predictable, even whilst the normative implications of the principles that might lead Kazimira to this place extend, for the most part, only to her experiences in the meeting room.

It is this final point, then, that captures the real tension between principled (revisionary) accounts of disagreement, epistemic self-trust, and the fundamental role that the latter has to play in our everyday intellectual lives. For, the possibility that following the relevant principles could produce such epistemological trauma is not a feature of this case specifically – but a result of the affective and psychological nature of *AST* and the fact that it is *AST* that plays the most significant role in our everyday intellectual lives. And given that nature, it may well be that the epistemic

consequences of following *any* such account will outstrip the normative implications of that account. And, as I see it, this possibility may constitute a significant problem for any account of disagreement of this kind.

Acknowledgements

Thanks to two anonymous reviewers for their thoughtful comments and suggestions on how to improve the paper, as well as the editors of the volume. Thanks also for discussion of an earlier version of this paper to the attendees and organisers of the conference *Disagreements: from Theory to Practice* at the University of Tartu, 2020, including especially Siobhan Kattago, Toomas Lott, Bruno Mölder, James Pearson, Gurpreet Rattan, and Kristin Schaupp. Section 3 of this paper reworks aspects of Chapter 5 of the author's PhD thesis: "Disagreement and deep disagreement: Should you trust yourself?", completed 2020, Sheffield University.

Address:

Simon Barker

Marati 4-19, Tallinn 11712

Estonia

Tel.: (00372) 58286190

Email: mail@simonbarker.is

References

- Alston, William P. (1986) "Epistemic circularity". *Philosophy and Phenomenological Research* 47, 1, 1–30.
- Burge, Tyler (2013) "Content preservation". In *Cognition through understanding: self-knowledge, interlocution, reasoning, reflection*. Oxford: Oxford University Press.
- Christensen, David (2009) "Disagreement as evidence: the epistemology of controversy". *Philosophy Compass* 4, 5, 756–767.
- Christensen, David (2011) "Disagreement, question-begging and epistemic self-criticism". *Philosophers' Imprint* 11, 6.
- Christensen, David (2016) "Disagreement, drugs, etc.: from accuracy to akrasia". *Episteme: A Journal of Individual and Social Epistemology* 13, 4, 397–422.
- Elga, Adam. (2007) "Reflection and disagreement". *Noûs* 41, 3, 478–502.
- Enoch, David (2010) "Not just a truthometer: taking oneself seriously (but not too seriously) in cases of peer disagreement". *Mind* 119, 476, 953–997.
- Faulkner, Paul (2011) *Knowledge on trust*. Oxford: Oxford University Press.
- Faulkner, Paul (2016) "Agency and disagreement". In Patrick J. Reider, ed. *Social epistemology and epistemic agency: decentralizing epistemic agency*, 75–90. London: Rowman and Littlefield International.

- Feldman, Richard (2006) "Epistemological puzzles about disagreement". In Stephen Hetherington, ed. *Epistemology futures*, 216–236. Oxford: Oxford University Press.
- Feldman, Richard (2007) "Reasonable religious disagreements". In Louise Antony, ed. *Philosophers without gods: meditations on atheism and the secular life*, 194–214. Oxford: Oxford University Press.
- Feldman, Richard (2009) "Evidentialism, higher-order evidence, and disagreement". *Episteme* 6, 3, 294–312.
- Foley, Richard (2001) *Intellectual trust in oneself and others*. Cambridge: Cambridge University Press.
- Fricker, Elizabeth (2016) "Doing (better) what comes naturally: zagzebski on rationality and Epistemic Self-trust". *Episteme: A Journal of Individual and Social Epistemology* 13, 2, 151–166.
- Hawthorne, John and Amia Srinivasan (2013) "Disagreement without transparency: some bleak thoughts". In David Christensen and Jennifer Lackey, eds. *The epistemology of disagreement: new essays*, 9–30. Oxford: Oxford University Press.
- Holton, Richard (1994) "Deciding to trust, coming to believe". *Australasian Journal of Philosophy* 72, 1, 63–76.
- Jones, Karen (1996) "Trust as an affective attitude". *Ethics* 107, 1, 4–25.
- Jones, Karen (2012a) "The politics of intellectual self-trust". *Social Epistemology* 26, 2, 237–251.
- Jones, Karen (2012b) "Response to Nancy Daukas". *Social Epistemology Review and Reply Collectiv*, 1, 10, 1–7.
- Kelly, Thomas (2010) "Peer disagreement and higher-order evidence". In Richard Feldman and Ted A. Warfield, eds. *Disagreement*, 111–175. Oxford: Oxford University Press.
- Kelly, Thomas (2013) "Disagreement and the burdens of judgment". In David Christensen and Jennifer Lackey, eds. *The epistemology of disagreement: new essays*, 31–53. Oxford: Oxford University Press.
- Lackey, Jennifer (2010) "A justificationist view of disagreement's epistemic significance". In Adrian Haddock et al., eds. *Social epistemology*, 298–325. Oxford: Oxford University Press.
- Lackey, Jennifer (2013) "Disagreement and belief dependence: why numbers matter". In David Christensen and Jennifer Lackey, eds. *The epistemology of disagreement: new essays*, 243–268. Oxford: Oxford University Press.
- Lasonen-Aarnio, Maria (2014) "Higher-order evidence and the limits of defeat". *Philosophy and Phenomenological Research* 88, 2, 314–345.
- Levy, Neil (2020) "The surprising truth about disagreement". *Acta Analytica* 1–21.
- Matheson, Jonathan (2009) "Conciliatory views of disagreement and higher-order evidence". *Episteme: A Journal of Individual and Social Epistemology* 6, 3, 269–279.
- Matheson, Jonathan (2015) *The epistemic significance of disagreement*. New York: Palgrave Macmillan.
- Peter, Fabienne (2019) "Epistemic self-trust and doxastic disagreements". *Erkenntnis* 84, 6, 1189–1205.
- Schafer, Karl (2015) "How common is peer disagreement? On self-trust and rational symmetry". *Philosophy and Phenomenological Research* 91, 1, 25–46.
- Tanesini, Alessandra (2020) "Virtuous and vicious intellectual self-trust". In Katherine Dormandy, ed. *Trust in epistemology*, 218–238. London: Routledge.
- Wedgwood, Ralph (2010) "The moral evil demons". In Richard Feldman and Ted A. Warfield, eds. *Disagreement*, 216–246. Oxford: Oxford University Press.
- Zagzebski, Linda Trinkaus (2012) *Epistemic authority: a theory of trust, authority, and autonomy in belief*. Oxford: Oxford University Press.