

Andrew Reisner

Draft of 30 June 2015

Peer Disagreement, Rational Requirements, and Evidence of Evidence as Evidence Against¹

[This is the penultimate or antepenultimate draft of the paper. Please do not cite without permission.]

O. INTRODUCTION

Worries about the significance of disagreement amongst actual, or perceived (perhaps with a good reason), peers has in recent years prompted a substantial body of work by epistemologists.² It is common in this literature to discuss a variety of normative issues that may arise from peer disagreement in terms of 'rationality'. This is regrettable. For roughly the last fifteen years, philosophers working on normativity have been growing ever more sensitive to the important differences between normative reasons, rational requirements, and requirements of reasoning. The term 'rationality', as used by epistemologists in this debate, can be used to make claims about any or any combination of normative reasons, rational requirements, and requirements of reasoning.³ 'Rationality', in the broader literature on theoretical reason, typically denotes requirements or permissions that appertain to relations amongst mental states, as opposed to considerations generally that count in favour of having particular beliefs or other mental states. Those fall under the rubric of 'normative reasons'.

Despite the terminological difficulties, it seems clear that the main focus of the peer

¹ I wish to express sincere my gratitude to Stewart Cohen, Michele Palmira, and Jessica Pepp for valuable discussions and feedback on this chapter.

² See De Ridder (2014) as an example of work on disagreement between idealised peers. For disagreement between agents who regard each other with reason as peers, see Enoch (2010).

³ Cases of ambiguous usage include Ballantyne and Coffman (2011), Kelly (2010) and White (2005). This is often the result of an uneasy slide from traditional talk of justification in epistemology to the use of reasons talk in concert with justification. For a brief discussion of the incorporation of reasons into epistemology, see Reisner and Steglich-Petersen (2011).

disagreement literature in epistemology is best understood as being on normative reasons for belief.⁴ Consequently, in this chapter I shall consider the possible import of peer disagreement with respect to rationality in its more regimented use and as distinct from the importance of peer disagreement with respect to the normative reasons it may provide or its import for the requirements of reasoning. My suggestion is that there is little upshot rationally to peer disagreement. A possible exception are cases in which peer disagreement is believed to be antecedently unlikely or in which it is assigned a low credence. At present, the question of the rational significance of peer disagreement has received little attention, and I hope to provide at least a starting point for a more thorough investigation of the topic.

The paper proceeds as follows. In §1, I introduce the distinction among normative reasons (just 'reasons' from here forward), rationality, and reasoning. In §2, I consider two views about the nature of evidence itself. In §3, I introduce an example of how evidence of evidence of p can be evidence against p . In §4, I discuss how this example limits the ways in which peer disagreement might affect what rationality requires of the respective peers. In §5, I discuss some ways in which improbable peer disagreements might affect what rationality requires of an agent. I offer some concluding thoughts about the limited importance of peer disagreement for theoretical rational requirements in §6.

1. REASONS, RATIONAL REQUIREMENTS, AND REASONING

Reading through the literature on peer disagreement in the early part of the now decade long boom in research on the topic reveals some surprises. At the start of this period, then recent work by John Broome, Derek Parfit, and Ralph Wedgwood had already begun to mark important

⁴ I thank Michele Palmira for suggesting this reading of the literature to me.

differences between normative reasons and the requirements of rationality.⁵ My own work on the distinction already dated back several years in the form of the beginnings of my doctoral thesis,⁶ although it would not appear in published form for some time to come.⁷ Niko Kolodny's now famous paper, 'Why Be Rational',⁸ which argues for the distinction, had been in wide circulation for not a short period prior to its 2005 publication in *Mind*. Attentive readers may have noticed a quite different pair of earnestists in work by Christopher Cherniak and Logi Gunnarsson.⁹ Yet this important distinction does not appear in a careful or systematic way in the early – and some of the later – literature on peer disagreement, even as interest in the relation and distinction between normativity and rationality has grown steadily.

As a consequence, before considering the importance of peer disagreement to rationality in its current more carefully specified sense, it will be useful to keep careful track of the respective domains for which peer disagreement may pose a puzzle. I shall discuss each of *reasons* and *oughts*, *rationality*, and *reasoning* briefly to explain the differences amongst them.

Oughts and Reasons

Ought is the final normative operator. In this way, it is distinguished from there being a *reason*, which is non-final normative operator. If one ought to believe x , then the matter is normatively settled, absent relevant changes in how things are. If one has a reason to believe x , it is still possible that one ought not to believe x . What one ought (or ought not) to believe will depend on whether and which other reasons bear on believing x .

⁵ Broome (1999), Parfit (2001), and Wedgwood (2003).

⁶ Reisner (2004).

⁷ Reisner (2011). The original paper that formed the basis of my doctoral work was 'Why Rational Requirements Are not Normative Requirements'. The main argument from this paper was folded into a book chapter, 'Is there Reason to Be Theoretically Rational?.'

⁸ Kolodny (2005).

⁹ Cherniak (1986) and Gunnarsson (2000).

There are competing views about the relationship between oughts and reasons. Some philosophers believe that reasons can be analysed in terms of oughts.¹⁰ Some believe that oughts may be analysed in terms of reasons.¹¹ My own view is that neither can be analysed in terms of the other. Whatever the relation in which they stand, they share an important feature: they are robustly normative.

It is difficult to give an illuminating account of what it is for something to be robustly normative.¹² Happily, an illuminating account is not necessary for present purposes. Robust normativity is the kind of normativity that is thought to be associated with moral requirements, the requirements of prudence, and epistemic requirements by philosophers who are realists – perhaps non-deflationary realists – about such requirements. It is different from the kind of normativity associated with systems or collections of rules, such as chess, etiquette, and doctoral thesis formatting.

Put another way, oughts and reasons, and any elaborations thereon (*prima facie* oughts, for example) are the currency of robust normativity. There are many other uses of 'ought' and 'reason' in English, and of their appropriate translations in other languages, but our concern is not with natural language usage.

Regrettably, remarks on the natural language usage of the word 'reason' from time to time exert a strong pull on theorising about reasons in particular. Would that it were not so. As it is, we must guard against a particular kind of quick move. This is the association, principally on grounds of linguistic commonality, of three conceptually distinct concepts that go by the name 'reason', such that they eventually merge into one.

¹⁰ See Broome (2013), Brunero (2013a) and Kearns and Star (2009) for a discussion of these analyses.

¹¹ This is the view expressed in Parfit (2001) and Skorupski (2011). However, they do not themselves call oughts 'oughts'. Nonetheless, it is perfectly clear that oughts are what they are discussing.

¹² Although his principal work on the subject is unpublished, Nicholas Shackel deserves the most credit for exploring different forces of normativity, distinguishing between 'correctness' and 'directivity'. See Shackel (MS).

The kind of reasons that are of interest here are *normative reasons*. A normative reason can be glossed as a consideration that counts in favour of something. A reason for belief is a consideration that counts in favour of believing something. There is a substantive issue as to what is required for one thing to count in favour of another, and indeed there is controversy even about the correct ontology for the *relata* in the counting-in-favour-of relation. The details of these controversies fortunately need not detain us here. We may understand the notion of counting in favour, and therefore of a normative reason, as what is being gestured at by the claim that seeing dark clouds counts in favour of believing that it will soon rain.

Separately, there is the notion of an operative or motivational reason. These are reasons that are thought to play some kind of psychological role for an agent, often in an actual or idealised reasoning process. Sometimes motivational or operative reasons are thought to be mental states themselves, other times something else that figures in *reasons for* explanations of belief or behaviour.¹³ However they are understood, they are conceptually distinct from normative reasons. This must be pointed out, because it is often claimed that to be a normative reason, something must be able to be an operative or motivating reason. I am sceptical that there is any interesting relation between normative and motivational reasons. If there is, however, it is one that must be explained, and it is not true in virtue of their being one and the same thing in a non-obvious way.

A final use of 'reason' is to mean 'explanation'. This use is generally not conflated with the normative use. Broome claims that normative reasons are weighing explanations of oughts. In this sense, he is a reductivist about normative reasons. Broome's view should be treated as a substantive metaphysical or conceptual claim that requires support.

¹³ I can make little sense of what philosophers are trying to get at in discussions of *reasons for*. Given that the role of claims about *reasons for* is normally to limit what can count as normative reasons, the two are clearly intended to be conceptually or metaphysically distinct.

Rationality

Much of the discussion of peer disagreement is presented in terms of 'rationality'. It can be difficult to tell what the sense of 'rationality' is that is being used. I shall introduce the current popular regimentation here. It is how I shall use 'rationality', and it would be sensible for philosophers writing on peer disagreement to adopt the now common convention, as it is a great aid to clarity.

It is sometimes said that the rational supervenes on the mental.¹⁴ It is more accurate to say that the requirements of rationality take as their object mental states, or relations amongst mental states. The supervenience formulation is introduced to avoid a confusion with reasons. What mental states it is rational for me to have depends primarily on what my other mental states are. One could imagine, as fitting attitudes theorists sometimes seem to, that all normative claims are fundamentally claims about what reasons there are for one to be in particular mental states. However, it is still open to say that the reasons are themselves exclusively, or primarily, non-mental facts. Thus, one could have reasons to be in mental state M , whilst one is rationally required not to be in mental state M .

An example will be useful. Take the rational requirement that if you believe p , then do not also believe not p . This requirement may be expressed in a wide-scope formulation:

$$\text{NCW: } \text{RR}(Bp \rightarrow B\neg p)$$

Or it may be expressed in a narrow-scope formulation:

$$\text{NCN: } Bp \rightarrow \text{RR-}B\neg p$$

¹⁴ Broome (2013), who follows Wedgwood (2003).

Whichever formulation is correct, each is thought to hold on account of a feature of an agent's mental life, namely that the contents of his beliefs are logically inconsistent, when he at once believes p and also believes not p . A typical rational requirement, whether wide or narrow in scope, has only mental states within the scope of the requirement, and in the case of narrow scope requirements, there are only mental states in the antecedents to the conditional.

What is being excluded are requirements of a form like this one: you are required to bring an umbrella, whenever it rains, even if you neither know nor believe that it is raining. I have deliberately been ambiguous about the scope of the requirement, because the point stands in both the wide-scope and narrow-scope versions; rationality does not include requirements like this one. This requirement is an ought, or perhaps a *prima facie* ought. It arises out of the reasons one has for not getting wet, when out in the rain.

The relation between rational requirements and reasons or oughts remains controversial. Fortunately, the correct answer to the debate about the normativity of rationality does not matter for the discussion in this paper, and thus it suffices to note that there is a distinction between reasons and oughts on the one hand and rationality on the other.¹⁵

Reasoning

Reasoning is a process. We commonly describe reasoning, or at least explicit reasoning, as a series of mental states. In the case of belief reasoning, both the states containing the premises and that which contains the conclusion are beliefs. When describing one's own theoretical reasoning, the mention of beliefs is typically omitted.¹⁶

¹⁵ In saying this, I assume that whether there are reasons to be rational, rationality is not reducible to reasons. I find the arguments in chapters 5 and 6 of Broome (2013) convincing on this point.

¹⁶ This notion of reasoning, specifically within the context of belief revision, appears in Harman (1984 and 1986), where

These are the easy things to say about reasoning. The rest is more difficult. To say that reasoning is a process is to say that it is more than just a sequence of disconnected propositional attitudes. Reasoning, as opposed to just having a series of attitudes, is distinguished by those attitudes with their contents having some sort of relation to each other. One arrives at the conclusion on account of the premises. Just how to understand *on account of* is unclear.

I shall not try to shed any light on the matter. There is something that makes the occurrence of a series of mental states in an agent's psychology an instance of reasoning rather than an instance of having a series of disconnected mental states, such that it is possible to have the same series of mental states on two different occasions and have one count as reasoning and the other not. What that something is remains controversial.¹⁷

In ordinary speech and in philosophy we make a further distinction between good reasoning and bad reasoning. It is important not to define reasoning in such a way that bad reasoning is impossible. As a regimentation for this paper, what normally falls under the heading of 'good reasoning' will either be what is required or permitted by reasoning. As with the requirements of rationality, I shall not assume that the requirements (or permissions) of reasoning are robustly normative. This leaves open the possibility that how one ought to reason may be at times contrary to the requirements of reasoning.

2. EVIDENCE AND EVIDENTIAL PARITY

In debates about epistemic peer disagreement, one needs a notion of epistemic peerdom. Typically, two necessary conditions are assigned to epistemic peers:

he helpfully distinguishes it from arguments that are used in reasoning.

¹⁷ This is the source of an interesting dialogue between Paul Boghossian and Broome. The outlines of the discussion may be found in Broome (2013). Boghossian's view is set out in detail in Boghossian (2012).

Evidential parity: Two individuals are epistemic peers with respect to p only if they share the same evidence with respect to p .

Shared capacities: Two individuals are epistemic peers with respect to p only if they are equally cognitively capable with respect to assessing whether p .

These conditions are not precise, and how they are made precise is significant for understanding what follows from the fact, or possibility, of epistemic disagreement. Some particularly interesting issues arise in the case of evidential parity.

The evidential parity condition is typically intended as a claim about agents' actual evidence, rather than what they take their evidence to be. Evidential parity will therefore look different depending on one's substantive theory of evidence. For illustrative purposes, I have chosen two extreme views about the nature of evidence, which I call 'strong evidential externalism' or just 'evidential externalism' and 'strong evidential internalism' or just 'evidential internalism'. Both are views that can be found in recent literature, although neither is especially popular in their strong forms. I use them as an illustrative example because they are easy to work with, and the relevantly external and internal characteristics of evidence on each view apply to the full spectrum of more nuanced theories of evidence.

Strong evidential externalism is a view that one finds primarily in the literature on theoretical reason. Although externalist views differ with respect to details, they typically share three features. The first is that the evidential relation is between facts (or true propositions) and propositions (either true or false). The second is that these facts need not be known or believed by agents for whom they are evidence. The third is that, although the agent need not know or believe the facts that constitute evidence for a particular proposition, there is some restriction on the accessibility of

the facts that is stricter than that which applies to those facts that constitute practical reasons.¹⁸

As a consequence, for two individuals to be evidential peers with respect to p , they need not share the same mental states with respect to p . They do, however, need to be similarly situated such that they have the same epistemic access, however defined, to the same facts about or pertaining to p .

By way of contrast, strong internalists about evidence think that an agent's beliefs, or credences, constitute his evidence. It is not always clear whether they think that the evidential relation is between doxastic states and propositions or doxastic states and other doxastic states, but the identity of the second *relatum* is not important in this context. We may note, then, that two individuals are evidential peers, according to strong internalism, just in case they share the same relevant doxastic states.¹⁹

One consequence of strong evidential externalism is that there are no rational requirements concerning first order evidence. That is because evidence is not a mental state, and rational requirements only govern relations among mental states. We cannot assume any principle of the sort that says: if there is a change in an agent's evidence, there is a change in what is rationally required of him.

Strong externalism about evidence does not preclude the possibility of there being second order rational requirements concerning evidence. One such requirement might be a version of *belief enkrasia*:

Belief enkrasia (BE): You are rationally required that if you believe that your total

¹⁸ Dancy (2002) and Skourpski (2011) offer versions of this form of strong externalism, and the position is perhaps implied in Parfit (2001).

¹⁹ With some modifications, notably from working with beliefs to working with credences, evidential internalism is the account of evidence in classical Bayesianism.

evidence supports proposition p to degree d , then you believe p .²⁰

BEF: $RR(BE_d p \rightarrow Bp)$.

The general plausibility of BE is much debated.²¹ I shall return to BE and its more general role in rational disagreement in §3.

Strong internalism about evidence is in a quite different position. The strong evidential internalist may accept that there are rational requirements on evidence, because evidence is itself a mental state.²² An example of such a requirement is the *direct evidential requirement*:

Direct evidential requirement (DE): You are rationally required that if your total evidence supports proposition p to degree d , then you believe p .²³

DEF: $RR(E_d p \rightarrow Bp)$.

The DE schema does not apply to all accounts of evidence. There is an asymmetry between BE and DE in this way; unlike DE, BE in principle could apply to any account of evidence. If externalist theories of evidence are correct, then there are no first order evidential requirements of rationality.

This point is significant for claims that there are rational requirements that arise from peer disagreement. If there are no DE-like rational requirements that apply to externalist accounts of evidence, then there are none to violate. Evidential externalism does not affect rationality through

²⁰ I have used all out belief in the consequent to avoid some complications arising from the use of credences.

²¹ See Brunero (2013b) and Reisner (2013) for discussions of its status as a rational requirement.

²² Note that on this way of classifying theories of evidence, so-called 'knowledge first' theories of evidence will come out as internalist. See Williamson (2000).

²³ It may be more natural to have a credence rather than a belief in the consequent of this conditional.

DE-like requirements.²⁴

What may be less obvious is that if externalism is true, then the fact of peer disagreement also does not bear on second order theoretical rational requirements like BE. According to the externalist account of evidence, my beliefs do not figure as direct evidence either for their own contents or for other propositions.²⁵ If you are my epistemic peer, then presumably your beliefs have the same epistemic status as my own, which is to say that they do not constitute direct evidence for the proposition about which we are in disagreement. If I (correctly, according to strong externalism) do not take your belief that not p as evidence for not p , then it is difficult to see why it would exert any pressure on my second order belief about whether there is sufficiently strong evidence for p .

Of course, I might believe that your belief is evidence that there is evidence that not p , perhaps because you are equally as good as I am at detecting external evidence. In this case, our disagreement is a sign that one of us erred either in appreciating the evidence we have or in reasoning with it. In the next section, I shall raise some doubts about whether this second order approach provides a way in which peer disagreement could change what we are rationally required to believe.

3. EVIDENCE OF EVIDENCE AS EVIDENCE AGAINST AND HIGHER ORDER EVIDENTIAL REQUIREMENTS

It is by now well known that evidence of evidence is not necessarily evidence, due to Brandon Fitelson's playing card example.²⁶ Things can be worse than mere failure: evidence of evidence can

²⁴ de Ridder (2014) offers an extensive discussion on important of evidential externalism for peer disagreement.

²⁵ This way of putting things is not meant to exclude my having a particular belief's standing as evidence for propositions concerning my belief states and other cases in which my having a particular belief or beliefs is an evidentially salient fact.

²⁶Fitelson (2012). This is the example. That the card is black is evidence that the card is the ace of spades. That the card is the ace of spades is evidence that the card is an ace. But, that the card is black is not evidence that it is an ace.

be evidence against. It may be useful to consider briefly an example in which that is the case.

The Polaroid murder

Sometime in the future, there are only two functioning Polaroid cameras left on earth. Each has a single remaining cartridge, which cannot be removed before it is used. Due to their scarcity, they have been rigged with a sophisticated biometric scanner that only allows them to be used by their respective owners, Paul and Brian. We should assume that these facts are known to all parties in the example.

A murder has been committed. The detective on the scene notices that there is what looks to him like a recently taken Polaroid lying on the body. He manages a quick but clear look at it as he approaches, but then the wind blows the picture away. If he saw the image correctly, it was a picture of the murderer's hand stabbing the victim. Given the angle and location of the body, only the murderer could have taken the photograph. Although neither have a (known) motive and neither are (known to be) violent, the detective knows that only Paul and Brian own a functioning Polaroid camera. To start the investigation, he pays Paul a visit.

The detective tells Paul that he is confident that there is evidence that Paul committed the murder and explains that he has likely seen a photograph that implicated either Paul or Brian. Paul pulls from his pocket a recently taken Polaroid photograph. It is not the one that the detective saw, but rather it appears to be a photograph of the photograph that the detective saw. Brian could have taken the second photograph and Paul could have stolen it, but the logistics are such that it is very much more likely that Paul would have taken the second photograph himself.²⁷

Let us call the proposition *that Paul committed the murder* 'M'. We can see that Paul's being in

²⁷ A simpler example of radical EE failure is this: Somebody said *x*; an anti-reliable person said *x*. For reasons that will become apparent, I prefer my example, because it does not rely on treating testimony as *prima facie* evidence. Stewart Cohen first drew my attention to the example in this footnote.

possession of the second photograph is evidence that he did not commit the murder, so it is evidence for $\neg M$. But the original photograph, for which the second is evidence, is itself evidence for M .²⁸ The second photograph is in the epistemologists' lingo a *rebutting defeater*. That is to say that it drives down the probability of a proposition by driving up the probability of its contrary. Importantly, however, there is still evidence that Paul committed the murder, it is just lower than it would have been, had Paul not been in possession of the second photograph. The second photograph is not an undermining defeater of the first.

4. PEER DISAGREEMENT, EVIDENCE OF EVIDENCE, AND RATIONAL REQUIREMENTS ON BELIEF

The lesson of the Polaroid Murder example is that peer disagreement will not change what one is rationally required to believe by serving as evidence that there is evidence for the contrary of one's own belief.²⁹ The general problem is that there is no free entailment from believing that there is evidence of evidence of p to believing that there is evidence of p . It is worth looking in more detail at why this is so with respect to particular rational requirements. We may return to two requirements that I have already mentioned: BEF and DEF. Let us consider BEF first:

$$\text{BEF: RR}(\text{BE}_i p \rightarrow \text{B}p)$$

BEF says that you are required that if you believe the total strength of the evidence for p is

²⁸ This assumes that the only evidence that Paul committed the murder was the photograph, and so the probability that Paul committed the murder is higher conditional on the detective's having seen the photograph (or on his believing that he saw it) than it is conditional on his not having seen it (or on his believing that he did not see the photograph). If one prefers, the example can be altered slightly so that Paul is more likely to have committed the murder than Brian.

²⁹ Christensen (2010) and Kelly (2010) seem to treat peer disagreement as an undermining rather than a rebutting defeater. This is significant for the construction of their arguments and others related to how, or whether, to revise one's credences in the face of peer disagreement. The argument in this chapter is not substantively changed by treating disagreement as an undermining defeater rather than a rebutting defeater.

sufficiently strong, then you believe p . BEF is likely false, for reasons I set out elsewhere,⁵⁰ but it is a good approximation of the kind of higher-order belief principle that features in discussions about rational disagreement.

Even if we are strong evidential internalists, the belief that there is peer disagreement has no direct bearing on this principle. BEF is just a rational requirement concerning beliefs about the sufficiency of one's total evidence. The belief that there is peer disagreement would have to connect to BEF indirectly, as it is itself not a belief about the sufficiency of one's total evidence.

The most obvious route to this is through what I shall call 'belief *enkrasia* aggregation':

Belief *enkrasia* aggregation (BEag): Where the strength of your total evidence (TE) constituting beliefs concerning p , $TE_{d_1}\{B\alpha...B\omega\}$, is sufficient for believing p , you are rationally required that if you believe that α is evidence for p to degree d_1 and you believe that β is evidence for p to degree d_2 ... and that ω is evidence for p to degree d_n , then believe that there is sufficient total evidence for p .

BEagF: $(TE_{d_1}\{B\alpha...B\omega\} \rightarrow BE_{s,p}) \rightarrow RR\{[(BE_{d_1}\alpha p) \& (BE_{d_2}\beta p) \dots \& (BE_{d_n}\omega p)] \rightarrow BE_t p\}$

BEag links your individual beliefs about the evidential support relations between individual pieces of evidence and p with your belief about the sufficiency of the total evidential support for p . The requirement should be read with enough latitude to treat the relations amongst the evidence and the total sufficiency as being non-monotonic, when appropriate.

BEag provides a place where changed beliefs about evidence can change the contents of one's rational requirements. Adding or subtracting a belief on the lefthand side of the conditional

⁵⁰ Reisner (2009 and 2014).

contained within the scope of the requirement, or changing the degree of evidential support for one or more of the beliefs on the lefthand side of the conditional contained within the scope of the requirement, in principle can change what belief, if any, belongs on the righthand side of the conditional contained within the scope of the requirement. Yet BEag does not apply to beliefs of the form: my peer believes not p , while I believe p . This belief, whether read *de dicto* or *de re*, does not have evidence as part of its content.

We need not consider other similar rational requirements. The problem with BEag for bringing beliefs about disagreement to bear on theoretical rational requirements is that the lefthand side of the conditional contained within the scope of the requirement only contains beliefs about evidence. We would thus need a rational requirement, the lefthand side of which does not contain beliefs about evidence and the righthand side of which does contain beliefs about evidence. The beliefs on the righthand side may be either about individual evidential relations or whether there is sufficient total evidence.

It is easy enough to produce a sample of such a requirement. We may call this one the 'interlevel belief requirement':³¹

Interlevel belief requirement (IB): Where your total evidence concerning p is $\{Bx, By, Bz\}$, and it is sufficiently strong for believing p , you are rationally required that if you believe x , believe y , and believe z , then you believe that there is sufficient evidence for p .³²

IBF: $(E_d\{Bx, By, Bz\} \rightarrow E_t p) \rightarrow RR[(Bx \& By \& Bz) \rightarrow BE_t p]$

³¹ I borrow the name, and the general idea, from Worsnip (MS).

³² This principle could be rewritten with credences rather than outright beliefs.

It contrasts with the direct or first order analogue of BEagF:

$$\text{DEagF: } (E_d\{Bx, By, Bz\} \rightarrow E_i p) \rightarrow \text{RR}[(Bx \& By \& Bz) \rightarrow Bp]$$

DEagF is just DEF with the individual pieces of internalist evidence, which are beliefs, written out individually and with the fact that they are beliefs made explicit. The beliefs Bx through Bz are just an example of a total set of beliefs that constitute sufficient evidence for p in DEagF. In this sense, we may treat DEagF as a fully unpacked version of DEF.

Rational requirements like DEagF are presumably explained as having the status of rational requirement on the grounds that they enforce a consistency relation on our doxastic states. In this case, the consistency relation is some kind of evidential or probabilistic consistency.

If we were to explain IBF in exactly the same way, we would be claiming that there is an evidential relation between the having of the beliefs in the antecedent of the conditional and the content of the belief in the consequent. The content of the belief in the consequent is a proposition about the strength of evidence for p .

To be precise, IBF relies on evidence for p being evidence of evidence for p . If evidence for p is not evidence of evidence for p , then it is difficult to see how IBF could be explained. So much the worse for IBF. For the same reason that evidence of evidence of p is not necessarily evidence for p , evidence for p is not necessarily evidence of evidence of p . If M is evidence of evidence for p and also evidence for not p , as in the Polaroid murder example, then it follows that M is evidence for not p and not evidence of evidence for not p . The interlevel belief requirement is not a rational requirement.

This leaves us with DEF. When I form the belief that I disagree with a peer, that belief could feature as direct evidence in one of two ways (on the assumption of strong evidential internalism), at least in principle. The first is that it could be direct evidence for a belief such as the belief that

Paul committed the murder.

It is difficult to explain on internalist accounts of evidence how this would work. My evidence consists of my beliefs or credences, or some restricted subset thereof. If we accept that peers have the same evidence, then we accept that they have the same set of doxastic states. The new information that comes from peer disagreement is that given the same doxastic states and the same degree of reasoning competence, my peer and I have arrived at different conclusions. Although I shall revisit this assumption presently, let us for the moment assume peer disagreement indicates that in some way I may have mishandled the evidence that I have.³³ This is evidence about my evidence. Given that we cannot assume that evidence of evidence for p is evidence for p , there is no simple way to make the fact of peer disagreement matter evidentially for an ordinary first order belief, such as that Paul committed the murder.

We have already noted that DEF is a rational requirement relating first order evidence in the form of beliefs (on the assumption of evidential internalism) to another belief. There is no restriction special restriction on the content of the belief in the condition's consequent. Thus it is possible to have a belief about the sufficiency of one's total evidence feature in DEF as the consequent. This is the second way in which my belief that I disagree with a peer could change what is rationally required of me. The belief that I disagree with a peer could feature in the antecedent of DEF, when the belief that there is sufficient evidence for p is in the consequent.

5. THE DIRECT EVIDENTIAL REQUIREMENT, ACTUAL PEER DISAGREEMENT, AND IMPROBABLE DISAGREEMENT

The strategy suggested at the end of the last section is to relate a belief about there being peer disagreement with respect to p to the belief that I have sufficient evidence for p . To do this, we

³³ At least, this is what is sometimes claimed. I think this claim is probably wrong on the merits, at least for any simple version of it, but I concede the point for the sake of argument.

would have to show that actual peer disagreement, or my belief that there is actual peer disagreement, is new evidence that I have missed some evidence concerning p and thus cannot assume that there is still sufficient evidence for p .

It is important to notice that I have said 'new evidence' in the preceding paragraph. It is normally the case that I shall take myself to have missed some evidence for p , when I am considering whether p . The importance of uncertainty about whether I have all the available evidence, if rationally significant at all to whether there is sufficient evidence for p , is already priced into the strength that I assign to the total evidence for p , provided that I am rational in the relevant way. The question, then, is whether discovering that there is actual peer disagreement raises or lowers the probability that I have missed, or mishandled, some of the evidence for p .

One way in which my belief that there is actual peer disagreement may raise the probability that I have missed or mishandled some evidence concerning p is by that disagreement being unexpected. To see why the probability that I have assigned to there being peer disagreement, prior to its occurrence, is significant for whether it is rational to believe I have missed or mishandled some evidence, it will help to look at the two cases. The first is of improbable disagreement, the second is of probable disagreement.

Probable and improbable disagreement

Electronic computers, although not flawless, are rather good at following instructions. People are less good at it, and even when doing our best, we quite often reason incorrectly. This makes certain kinds of conversations between two or more computers and two or more humans quite different in character, especially when the computers disagree amongst themselves and the humans disagree amongst themselves.

We can think of two or more computers as being epistemic peers, or the appropriate analogue of that if you do not like the idea of computers having an epistemic status, just in case they meet

three conditions. 1) They have the same hardware. 2) They have the same software. 3) They are in the same initial information state and receive the same new instructions and data in the same sequence. If we make the further assumption that the software and hardware are both well implemented, computer epistemic peers will almost never disagree with each other.⁵⁴

The extraordinary reliability of computers at executing their programmes, at least under favourable conditions, is such that the chance that two computers will make an error at the same time for the same operations is minuscule. It is this fact that led NASA to use a voting system for its computer systems used on the space shuttle fleet.⁵⁵ Three computers simultaneously performed critical calculations during the space shuttle missions. Because disagreement between peer computers was highly improbable, if two of the three computers agreed, the outlying computer's results were discarded.⁵⁶

We can describe an alternative scenario in which a space shuttle only had two computers. In this case, disagreement between the computers would be regarded with great concern. Absent evidence about which computer made the mistake, mission control would no doubt recommend some sort of recalculation procedure, perhaps after checking the computers for damaged hardware or corrupted software.

If we imagine a human agent with computer-like reliability and accuracy disagreeing with a peer about a matter that was computationally tractable, he would be in the same position as mission control would be in with respect to the two computer disagreement case. That there is disagreement indicates that one of the two peers has mishandled the evidence. Assuming that

⁵⁴ In fact, we will need some further assumptions about the set-up of the computers. Two important assumptions are that their programmes are computationally tractable and that the programmes do not employ randomisers.

⁵⁵ It is well worth looking at Tomayko (1988), which records the history of NASA's experience with computers in Spaceflight. I thank Paul Niles for directing me to this resource and for numerous valuable discussions about this and other topics.

⁵⁶ Assuming that one computer is not undergoing some kind of systematic failure.

there is no further evidence as to which peer has done so, one peer's belief that there is actual peer disagreement stands as new evidence that he has missed or mishandled his evidence.

Probable evidence

In other cases, when disagreement is, or is thought to be, very likely, the fact of actual disagreement has little impact on the probability that we have missed, or mishandled, our evidence. As the probability of error increases, it becomes increasingly surprising that two peers agree rather than that they do not. Peer disagreement ratifies the prior expectation in cases where non-systematic error is likely.³⁷ Indeed, in some cases it may have no impact at all.³⁸

Examples of this situation abound in philosophy. I know that it is highly likely that I disagree with someone whom I regard as an epistemic peer about a variety of substantial philosophical positions that I currently believe. I may or may not have updated my philosophical beliefs appropriately, given the high probability that I assign to peer disagreement. However, discovering that there is an actual peer disagreement, when I have already assigned a high probability to one's occurring, does not raise by very much the probability that I have erred in my reasoning.

A rational requirement?

The preceding discussion suggests that the rational importance of actual peer disagreement, if it has any at all, may depend on the antecedent probability assigned to the proposition that there will be peer disagreement. It has done so without having given a concrete rational requirement for linking beliefs about actual peer disagreement to beliefs about the sufficiency of an agent's

³⁷ Systematic errors should be understood as those that are likely to produce common judgements by those who share in them, as opposed to non-systematic errors, which should be assumed to produce common judgements infrequently.

³⁸ Predicates of taste may be one such case. According to relativists, disagreement is expected and in many circumstances devoid of normative or rational significance. See MacFarlane (2014) for a complete discussion of the current state of the art.

evidence.

Before giving the precise rational requirement, I shall offer the *actual disagreement schema*:

Actual disagreement schema (ADS): The degree to which actual peer disagreement about p constitutes evidence that a peer's credence towards p is incorrect partially depends on the credence that that peer has antecedently assigned to the probability that there will be peer disagreement about p . The relation is such that actual disagreement about p counts for more when the antecedent probability assigned to there being disagreement is low than when it is high.

ADS is a way of giving shape to the thought that actual peer disagreement can sometimes make us less confident that we have got the right conclusion, because we may have missed or mishandled our evidence.

What ADS does not do is tell us how to complete the story of how believed actual peer disagreement will change our rational requirements concerning what to believe about the proposition that is the subject of the disagreement. It may be possible to formulate a rational requirement in which there is a link between uncertainty about the state of one's total evidence for p and whether there is sufficient evidence to believe p . If this can be done, and if we accept the higher order belief requirement, belief *enkrasia*,⁵⁹ then we may establish some rational link between believing that there is actual peer disagreement about p and whether to believe p .

⁵⁹ You are rationally required that if you believe that your total evidence supports proposition p to degree d , then you believe p .

6. CONCLUSION

This chapter has been about a narrow question concerning peer disagreement: namely whether peer disagreement affects an agent's theoretical rational requirements concerning beliefs about the proposition that is the subject of the disagreement. This is a separate question from that of whether it changes an agent's reasons for belief or affects how an agent ought to reason about a particular proposition. The debate on epistemic peer disagreement would benefit from a greater regimentation of the use of the term 'rational'. At present, it is unclear whether authors are using the term to refer to normative reasons, requirements of reasoning, or requirements of rationality in the more regimented sense of 'rationality'.

The contention in this chapter is that disagreement between epistemic peers does not clearly have an influence on what an agent is rationally required to believe, or on which collections of beliefs he is rationally required to have. This is so mainly on account of two considerations. The first is that evidence of evidence of p can be evidence against p . A natural belief to form on account of peer disagreement that is that the disagreement constitutes evidence that there is evidence for the contrary of one's belief. Because this does not entail that there is evidence for the contrary of one's belief, or that there is some other sort of evidence against one's belief, there is no direct rational upshot for one's first order belief.

The second consideration is that rational requirements govern mental states. Whatever the evidential situation is when there is peer disagreement, it only has a bearing on one's rational requirements when there is a relevant belief associated with it. The mere fact that there is disagreement is not relevant to one's rational requirement. That one believes that there is disagreement could in principle be important. The difficulty is in explaining how.

I have suggested that the belief that there is actual peer disagreement may play a role in determining what rationality requires of us, when we also believe that peer disagreement about a particular proposition is improbable. However, it remains to be seen whether this is really so. One

would need to accept an internalist theory of evidence and to defend both something like belief *enkrasia* and also a principle linking uncertainty about the state of one's evidence for p with beliefs about the sufficiency of one's evidence for p . I have not attempted either task in this paper. Until they are both discharged, it remains unclear as to whether there are any rational upshots to believing that there is actual peer disagreement.

Works Cited

- Ballantyne, N. & Coffmann, E. J. (2011). Uniqueness, Evidence, and Rationality. *Philosophers Imprint* 11: 1-13
- Boghossian, P. (2012). What Is Inference? *Philosophical Studies* 1: 1-18.
- Broome, J. (1999). Normative Requirements. *Ratio* 12: 398-419.
- Broome, J. (2013). *Rationality through Reasoning*. London: Wiley Blackwell.
- Brunero, J. (2013a). Reasons as Explanations. *Philosophical Studies* 165: 805-824.
- Bruneo, J. (2013b). Rational Akrasia. *Organon F* 4: 546-566.
- Cherniak, C. (1986). *Minimal Rationality*. Cambridge: MIT Press.
- Christensen, D. (2010). Higher-Order Evidence. *Philosophy and Phenomenological Research* 81: 185-215.
- Dancy, J. (2002). *Practical Reality*. Oxford: Oxford University Press.
- de Ridder, J. (2014). Why Only Externalists Can Be Steadfast. *Erkenntnis* 79: 185-199.
- Fitelson, B. (2012). Why Evidence of Evidence Is not (Necessarily) Evidence. *Analysis* 72: 85-88.
- Gunnarsson, L. (2000). *Making Moral Sense: Beyond Habermas and Gauthier*. Cambridge: Cambridge University Press.
- Harman, G. (1986). *Change in View*. Cambridge: MIT Press.
- Harman, G. (1984). Logic and Reasoning. *Synthese* 60: 107-127.
- Kearns, S. & Star, D (2009). Reasons as Evidence. *Oxford Studies in Metaethics* 4: 215-242.
- Kelly, T. (2010). Peer Disagreement and Higher Order Evidence. In A.I. Goldman and D. Whitcomb, (eds.) *Social Epistemology: Essential Readings*. Oxford: Oxford University Press: 183-217.
- Kolodny, N. (2005). Why Be Rational? *Mind* 114: 509-563.
- MacFarlane, J. (2014) *Assessment Sensitivity: Relative Truth and its Applications*. Oxford: Oxford University Press.
- Parfit, D. (2001). Reasons and rationality. In D. Egonsson, J. Josefsson, B. Petersson, and T. Rønnow-Rasmussen (eds.) *Exploring Practical Rationality*. Aldershot: Ashgate, 17-39.
- Reisner, A. (2004). *Conflicts of Normativity*. DPhil Thesis. University of Oxford.

- Reisner, A. (2009). The Possibility of Pragmatic Reasons for Belief and the Wrong Kind of Reasons Problem. *Philosophical Studies* 145: 157-172.
- Reisner, A. (2014). A Short Refutation of Strict Normative Evidentialism. *Inquiry* 5: 1-9.
- Reisner, A. (2013). Is the *Enkratic* Principle a Requirement of Rationality? *Organon F* 20: 436-462.
- Reisner, A. & Steglich-Petersen, A. (2011) Introduction. In A. Reisner and A. Steglich-Petersen (eds.) *Reasons for Belief*. Cambridge: Cambridge University Press: 1-10.
- Shackel, N. (MS). Two Kinds of Normativity.
- Skorupski, J. (2011). *The Domain of Reasons*. Oxford: Oxford University Press.
- Tomayko, J.E. (1988). The History of Computers in Spaceflight: The NASA Experience. NASA Contractor Report 182505. Available at <http://history.nasa.gov/computers/CompSPACE.html>.
- Wedgwood, R. (2003). Choosing Rationally and Choosing Correctly. In S. Stroud and C. Tappolet, (eds.) *Weakness of Will and Practical Irrationality*. Oxford: Oxford University Press: 201-229.
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.
- White, R. (2005). Epistemic Permissiveness. *Philosophical Perspectives* 19: 445-459.
- Worsnip (MS). The Conflict of Evidence and Coherence.