

Dilemmas, Disagreement, and Dualism

Liz Jackson, Ryerson University (Lizjackson111@gmail.com)

For *Epistemic Dilemmas: New Arguments, New Angles* (Routledge)

Kevin McCain, Scott Stapleford, and Matthias Steup, eds

Abstract: This paper introduces and motivates a solution to a dilemma from peer disagreement. Following Buchak (forthcoming), I argue that peer disagreement puts us in an epistemic dilemma: there is reason to think that our opinions should both change and not change when we encounter disagreement with epistemic peers. I argue that we can solve this dilemma by changing our credences, but not our beliefs, in response to disagreement. I explain how my view solves the dilemma in question, and then offer two additional arguments for it: one related to contents and attitudes, and another related to epistemic peerhood.

Keywords: epistemic dilemma; disagreement; conciliationism; steadfastness; belief; credence; belief-credence dualism; epistemic peer; content; attitude

1. Introduction

What should we do when someone who is smart and well-informed disagrees with us? Should we change our opinion, or hold fast to our previous viewpoint? This question has divided epistemologists, and those working in the epistemology of disagreement have spilled much ink defending various answers.¹

One concept that's been central to the disagreement debate is an *epistemic peer*. An epistemic peer is roughly someone who is epistemically "on a par" with you—they have similar evidence to you, and are approximately equally reliable, epistemically virtuous, and free from bias about the matter in question. While later, I'll argue that epistemic peerhood need not play a central role in the epistemology of disagreement, it's important for understanding the debate thus far. For now, we'll understand the central question of peer disagreement as: should one alter their opinion when they knowingly encounter a disagreeing epistemic peer?

In this paper, following Buchak (forthcoming), I argue that this central question of peer disagreement presents us with a dilemma; we have good reason to answer both 'yes' and 'no' to this question. I then offer and motivate a particular solution to that dilemma. This solution revolves around what I call *belief-credence dualism*, the view that we have two attitudes—beliefs and credences—and neither reduces to the other. In particular, I'll argue that we should alter our credences, but not our beliefs, when encountering an epistemic peer.

This paper proceeds as follows. In **Section 2**, I explain how peer disagreement leads to an epistemic dilemma. That is, there's good reason to think that we should both revise and not revise our opinions in response to peer disagreement. In **Section 3**, I offer a solution to this dilemma: if belief-credence dualism is true, this allows for the possibility that we should modify our credences, but not our beliefs, in response to peer disagreement. In **Sections 4 & 5**, I motivate this solution; my first motivation relates to the nature of contents and attitudes, and the second relates to epistemic peerhood. In **Section 6**, I address some worries about the view I've proposed, and in **Section 7**, I conclude.

¹ Some key works in the epistemology of disagreement include Kelly (2005), Christensen (2009), Lackey (2008), and the essays in Warfield and Feldman (2010) and in Christensen and Lackey (2013), and Pittard (2019).

2. The Dilemma from Disagreement

Buchak (forthcoming) argues that the disagreement debate lends support to propositions that seem to conflict with each other. Here we will focus on two, where ‘should’ indicates a requirement of epistemic rationality. The first is:

- (1) Our opinions should change when we encounter peer disagreement.

(1) reflects the intuition behind views that have come to be known as *conciliationist* (Christensen 2007, 2009, 2016; Elga 2007; Feldman 2007; Kornblith 2010; Turnbull and Sampson 2020; Fleisher forthcoming, among others). The thought is that, if you encounter someone that roughly shares your evidence, who has similar reliability and epistemic virtues, then it seems dogmatic and closeminded to simply ignore the fact they disagree with you by not altering your opinion—and in most cases, moving your opinion closer to your peer’s opinion.² Consider the classic restaurant case from Christensen (2007: 193) that is often used to motivate conciliationism:

Suppose that five of us go out to dinner. It’s time to pay the check, so the question we’re interested in is how much we each owe. We can all see the bill total clearly, we all agree to give a 20 percent tip, and we further agree to split the whole cost evenly, not worrying over who asked for imported water, or skipped desert, or drank more of the wine. I do the math in my head and become highly confident that our shares are \$43 each. Meanwhile, my friend does the math in her head and becomes highly confident that our shares are \$45 each. How should I react, upon learning of her belief?

Assuming you and your friend are epistemic peers in this case, Christensen concludes that it is “obvious” that you should change your opinion about how much you each owe after encountering your friend’s disagreement. Conciliationists are divided on how exactly we should alter our opinions—some endorse an “equal weight” view on which we should give our peer’s opinion the same weight as our own opinion (see Elga 2007); others discuss more nuanced updating rules (Easwaran et al 2016). Nonetheless, they agree that peer disagreement ought to change our opinions.

Others in the epistemology of disagreement argue that (1) is false:

- (2) Our opinions should *not* change when we encounter peer disagreement.

(2) is a version of what has become known as the *steadfast* view. As steadfasters point out, you seem to lose something epistemically valuable if we constantly defer to the opinions of others (van Inwagen 1996; Rosen 2001; Kelly 2005, 2007, 2010, 2013; Pettit 2006; Conee 2010). If conciliationism is true, virtually no one can rationally have a strong opinion about controversial matters. Some have also argued that always changing your views in response to disagreement leads to a problematic spinelessness—a requirement to give up our most deeply-held beliefs (see Elga 2007: 494; Sherman 2015; Fritz 2018). Not only does conciliationism seem to potentially lead to problematic skepticism, but there’s also evidence that remaining steadfast in the face of disagreement has epistemic benefits. For instance, Kitcher (1990, 1993) and Muldoon (2013) argue that diverse opinions among researchers makes it more likely that a group will reach the truth in the long run. Loughheed (2020) directly frames these facts as a challenge for conciliationism.

² Christensen’s doctor case is one potential exception to the rule that conciliationism always requires us to move closer to our peers’ opinions (2009: 759).

One way to think about what is going on here—the tension between (1) and (2)—is that peer disagreement puts us in an *epistemic dilemma* (Conee 1994). On the one hand, it seems dogmatic to hold onto your opinions when smart people disagree, but on the other, changing your opinion simply because someone disagrees seems spineless and overly deferent. Thus, there's pressure to both change and not change our views in response to disagreement. But *prima facie* we can't do both; (1) and (2) appear to be incompatible.

Before proceeding, note that some—most notably, King (2012)—have argued that peer disagreement rarely occurs, because it's seldom true that two people actually share the exact evidence, epistemic virtues, reliability, etc. One might reason from King's argument: since we never encounter epistemic peers, we simply don't face the problem of how to change our opinions in response to peers, so it doesn't matter if (1) and (2) conflict. Two things to note on this issue. First, as Kelly (2005) and Ballantyne (2014) note, the problem of peer disagreement seems to arise even with merely *potential* disagreeers—or, to use Ballantyne's phrase, “counterfactual philosophers”. So even if there are no actual peers, the conflict between (1) and (2) is still a problem: for you and a possible peer. Second—and more importantly—it doesn't matter if we have perfect or exact epistemic peers. Suppose that conciliationism is true. We should still change our opinions in response to others, even if they are slightly more reliable or slightly less reliable than we are—maybe we should change a bit more in the former case, and a bit less in the latter case. I'll address this extensively in Section 5 below, and show how my solution enables us to respond more flexibly to disagreement. Generally, the problem of disagreement doesn't depend on the existence of perfect epistemic peers.

3. The Dualist Solution

We've seen that there's good reason to accept both:

- (1) Our opinions should change when we encounter peer disagreement.
- (2) Our opinions should *not* change when we encounter peer disagreement.

But these seem to conflict; it's not clear that they can both be true. In this section, I'll argue that we can affirm versions of both (1) and (2), because “opinion” is ambiguous between two mental states: belief and credence.

Belief is a familiar attitude; belief is taking something to be the case or regarding it as true (Schwitzgebel 2019). I believe that modus ponens is valid and that my coffee has gotten cold. According to the tripartite model, there are three doxastic attitudes one can take toward a proposition *p*: believe *p*, disbelieve *p*, and withhold belief (where withholding means being undecided on whether *p*).

But consider: I believe both that modus ponens is valid and that my coffee has gotten cold, but my attitude toward these propositions isn't exactly the same—the former is more probable. To capture this, epistemologists appeal to another propositional attitude, called *credence*. The term ‘credence’ was originally introduced to mean something like subjective probability (Ramsey 1926: 166ff, Jeffrey 1965, de Finetti 1974, Eriksson and Hájek 2007), but many epistemologists now closely connect credence to our everyday notion of confidence (Schubach 2018: 191, Moon 2019: 276–7). Credences are more fine-grained than beliefs and are often given a value on the $[0,1]$ interval, where 1 represents maximal confidence that *p* is true, and 0 represents maximal confidence that *p* is false. For example, my credence that modus ponens is valid is very close to 1,

but my credence that my untouched coffee is cold is around 0.8. Unlike belief, there are (in principle) an infinite number of credences one could take toward a proposition.

There are several views about how belief and credence relate to each other (see Jackson 2020-b for an overview). We will focus on one: *belief–credence dualism*. On this view, belief and credence are independent mental states—we have both, and neither reduces to the other. While there are many arguments for dualism (see Ross and Schroeder 2012; Buchak 2014; Staffel 2017; Jackson 2019-b; Weisberg 2020), especially influential is the idea that belief and credence are two “tools” in our mental toolbox. Sometimes, it is useful to see the world in an on-off way—to simply affirm, deny, or remain neutral on propositions. Believing *p*, for example, allows us to take a stand on the truth of *p*. When one believes *p*, one represents the world as if *p* is true. Ross and Schroeder (2012) argue that high credences cannot play this same role. Since a 0.9 credence in *p* treats the possibility of not-*p* as live, high credences don’t let us take a stand or have a view of the world. Nonetheless, credences play another role—tracking one’s exact level of evidential support for a proposition. In some circumstances, we need a more fine-grained picture of the world than belief-attitudes allow—for example, when buying car insurance, it’s useful not to just acknowledge that you withhold belief on whether you’ll get in a wreck next year, but to think more precisely about the exact chance of that happening. In other words, you should consider your credence, not just your belief-attitude. In the same way that, when painting, we use both a roller brush and a detail brush to balance efficiency and accuracy, when forming attitudes and reasoning, we have both beliefs and credences; both have an important role to play in reasoning and acting.³

How does this relate to peer disagreement? If we have both beliefs and credences, this opens up additional possibilities for how we should respond to disagreement (Jackson 2019-a). More specifically, when we encounter peer disagreement about *p*, here are some possibilities for what we should change:

- (i) *Both* our credence in *p* and our belief-attitude in *p*.
- (ii) *Neither* our credence in *p* nor our belief-attitude in *p*.
- (iii) Our belief-attitude in *p*, but not our credence in *p*.
- (iv) Our credence in *p*, but not our belief-attitude in *p*.

Option (i) requires conciliating with both attitudes, and option (ii) requires remaining steadfast with both attitudes. Given these, the conflict between (1) and (2) still arises. However, on options (iii) and (iv), we can maintain that both (1) and (2) are true—at least in some sense. Here, I’ll argue for option (iv)—and in the next two sections, I’ll explain why we should prefer it to (iii). More precisely, I’ll argue that we can replace (1) and (2) with:

- (1*) Our credences should change when we encounter peer disagreement.
- (2*) Our beliefs should *not* change when we encounter peer disagreement.

In virtue of (1*), this response to disagreement doesn’t require us to be dogmatic or closeminded. We can acknowledge and give weight to our peer’s dissenting opinion. In virtue of (2*), however, this view doesn’t result in spinelessness or a loss of epistemic benefits. There’s no requirement to give up our most deeply-held beliefs, and epistemic communities can still glean the benefits of

³ Arguably, empirical evidence is also relevant to whether belief–credence dualism is true. For an excellent summary of the empirical support for dualism, see Weisberg (2020). For views on how belief and credence might affect action, especially in a dualist framework, see Ross and Schroder (2012), Buchak (2014), and Jackson (2019-b).

diverse perspectives. The fact that (1*) and (2*) enable us to avoid the dilemma from disagreement is an initial motivation for them.⁴

Note also that (1*) and (2*) also offer a solution to a widely-acknowledged problem for conciliationism: the self-undermining problem. The problem is that if conciliationism is true, it doesn't seem like we can rationally believe it, since many smart people disagree about conciliationism (see Christensen 2009: 762; Sampson 2019; Fleisher forthcoming). But this is only a problem if one holds to belief-conciliationism. According to (1*) and (2*), conciliationists should have a lower credence in, but can nonetheless believe, conciliationism. They can thus have no qualms about rationally believing their favorite theory, even though they ought to conciliate with peers by lowering their credences in it. Now, we'll consider two additional arguments for (1*) and (2*).

4. Contents and Attitudes

We've seen that conciliationism about credence and steadfastness about belief provides a potential way out of the dilemma from disagreement. A second reason to prefer this view has to do with the mental structure of belief and credence—specifically, the distinction between *contents* and *attitudes*. On the traditional model of belief, for every proposition one has considered, there are only three belief-like attitudes one can take to it—believe *p*, withhold on *p*, or disbelieve *p*. Because there are only three attitudes, it's hard to see what conciliationism should look like in a belief framework. If you believe *p* and I disbelieve *p*, and we are peers, maybe we should both withhold belief on *p*. But suppose you believe *p*, and I withhold on *p*? Or you withhold on *p*, and I disbelieve *p*? It's not at all clear how we should modify our beliefs to conciliate in these cases (see Jackson 2019-a).

One might suggest that, if I believe *p* and you withhold, maybe we should conciliate by both believing probably-*p*. Relatedly, one might think that we should conciliate by changing our attitude toward a higher-order proposition, such as “the evidence that my peer and I share supports *p*.”⁵ These suggestions have intuitive appeal. However, neither tells me what attitude I ought to take toward the bare proposition, *p*. In this case, *p* is a proposition I've entertained, and, since conciliationism is true and I've encountered you, a disagreeing peer, I shouldn't believe it. But I also shouldn't withhold, as that would amount to over-conciliating by simply adopting your attitude. I don't have reason to think that you're *more* reliable than me. And surely I shouldn't disbelieve *p*. Belief-conciliationism actually appears to lead to its own epistemic dilemma, of the exact sort that Conee (1994) discusses: *S* is in an epistemic dilemma if “*S*'s epistemic reasons argue that *S* should not believe *p*, should not deny *p*, and should not withhold judgment on *p*...none of the three doxastic attitudes *S* might take toward *p*—belief, denial, and withholding judgment—is obviously reasonable” (475). The belief framework is thus simply too coarse-grained to handle even basic cases of conciliation.

⁴ It's worth comparing (1*) and (2*) to the view of disagreement recently defended by Buchak (forthcoming). First, rather than focusing on belief and credence, Buchak focuses on one's ‘opinion’ and ‘one's assessment of the evidence.’ Further, while she notes in footnote 22 that her view compliments a view like mine—and several of her arguments may also support the view I'm defending—she's nonetheless concerned with views that hold belief and credence to be closely related and maintain a dependency link between the two attitudes. Here, I'm assuming that belief and credence are descriptively independent, as dualism suggests, and will later note implications my view has for the relationship between rational belief and rational credence (i.e. for normative independence).

⁵ Thanks to Kevin McCain.

One might reply that if you believe p and if I withhold on p , we don't actually disagree. Since withholding belief is an "undecided" attitude, it doesn't amount to taking a stance on anything. It's not apt to count as disagreement if one party takes a stand and the other doesn't. So the only cases of disagreement we need to worry about are cases where you believe p and I disbelieve p .⁶

In response, note that there are two reasons I might withhold belief on p . On the first, I withhold belief because I haven't thought much about p or examined much evidence; maybe I hadn't considered whether p until just now. On this conception, you might not conciliate with me, but this isn't because we don't disagree. This is because you have good reason to think I'm not an epistemic peer on the matter in question.

The second reason I might withhold belief on something is because I've thought long and hard about the matter, I've considered the evidence, and I'm convinced withholding is the correct response. You, my peer, believe p . In this case, you can't easily discount me as a peer. If belief-conciliationism is true, it appears that we should both modify our attitudes toward p in some way. But again, it's not at all obvious what modifications we should make. Thus, we should look beyond belief to capture the intuitions behind conciliationism.

Credences, on the other hand, are significantly more fine-grained than beliefs. Since one can adopt a credence toward p anywhere on the interval from $[0,1]$, there are (at least in principle) an infinite number of credences one can take toward a proposition. This allows for a lot more flexibility when conciliating with a peer. Let's consider again the example above, on which you believe p and I withhold on p . Suppose that your credence in p is 0.9, and mine is 0.5. On a simple picture on which you should conciliate by splitting the difference with your peer (i.e. the *split-the-difference* view), if credal conciliationism is true, we should both adopt a credence of 0.7 toward p . Or, if you're at 0.1 and I'm at 0.7, we should both adopt a 0.4 credence. Credences allow for a flexibility in conciliating that beliefs do not.

Of course, this doesn't mean that, if credal conciliationism is true, we should always simply split the difference with our peer's credences. Christensen (2009: 759) provides the following case:

I am a doctor determining what dosage of a drug to give my patient. I'm initially inclined to be very confident in my conclusion, but knowing my own fallibility in calculation, I pull back a bit, say, to 0.97. I also decide to ask my equally qualified colleague for an independent opinion. I do so in the Conciliatory spirit of using her reasoning as a check on my own. Now suppose I find out that she has arrived – presumably in a way that also takes into account her fallibility – at 0.96 credence in the same dosage.

You are technically encountering peer disagreement, as your colleague is less confident in the proposition than you. According to the split-the-difference view, you should be slightly less confident that you've found the correct dose, and alter your credence to 0.965. However, Christensen (2009: 759) concludes that "it seems that the rational thing to do is for me to increase my confidence that this is the correct dosage, not decrease it as difference-splitting would require." Even though your colleague is less confident, she is confident enough that it confirms, rather than calls into question, your conclusion about the dosage.

Of course, not all will share Christensen's intuitions about this specific case, but there are additional problems with the split-the-difference view. Others (e.g. Gardiner 2014; Russel et al

⁶ Thanks to Jaakko Hirvelä.

2015; Easwaran et al 2016) have argued that mechanical difference splitting leads to failures of commutativity—cases in which the order in which one encounters peers affects their final credences. For example, if I'm at 0.9 in p and encounter peer 1 at 0.5 in p , move to 0.7, then encounter peer 2 at 0.1, I end up at 0.4. I'll come to a different credence than if I encounter the peer 2, at 0.1, first—in this case, I will immediately move to 0.5, and then find that peer 1 agrees with me.

Thus, simple difference-splitting may not be the correct rule for all credal conciliation—even if it's a useful heuristic for certain cases. Thankfully, there are many other possible rules we might utilize, including both conditionalization and the “UPCO” rule that Easwaran et al (2016) propose. Conditionalization, or simply treating the fact that another has a particular credence as a piece of evidence then updating your credences on that evidence, gives accurate verdicts, but is very computationally demanding. Easwaran et al propose UPCO as a useful heuristic that mimics conditionalization in most circumstances. Here, we need not take a stand on the best rule to use in every circumstance, or how we should trade off simplicity and accuracy when conciliating. The main point is that it's natural to think that, if we are going to conciliate with one attitude, it should be credence, rather than belief. It is not clear how we'd conciliate in a coarse belief framework, and credence gives us the flexibility to accommodate different rules of revision.

5. Epistemic Peerhood

Recall that epistemic peers are people who are our epistemic equals. Our peers have similar (or the same) evidence, reliability, and epistemic virtues. The concept of an epistemic peer has been relatively central to the disagreement debate, and many frame the debate in terms of what we should do when we encounter epistemic peers.

However, as King (2012) contends, peer disagreement is very rare. King argues that peer disagreement minimally requires four things: one, someone with the same evidence, two, who is roughly equally reliable at evaluating that evidence, three, who disagrees with you on something (i.e. it's not just a verbal dispute), and four, you need a good reason to think that these first three conditions are met (252–3). But finding someone who merely has the same evidence as you is difficult—we come to the table with different experiences and background beliefs, and these are often difficult to share. Further, most people aren't at your *exact* level of reliability—many are probably slightly better or slightly worse at evaluating the evidence. If peerhood requires all four of these conditions being met at the same time, it seems like true peer disagreement rarely, if ever, occurs.

But note that the lack of peer disagreement doesn't mean conciliationism is never appropriate. Suppose conciliationism is true. If we disagree, but you are slightly my epistemic superior (e.g. you're slightly more reliable or have slightly better evidence), rather than my peer, then I should still conciliate with you (and maybe slightly more than I would have otherwise). If you are my epistemic inferior, but only by a little, then I should still conciliate (but maybe slightly less). If you're an expert, maybe I should defer to you, and if you're a total novice, maybe I shouldn't conciliate at all.

Let's consider an example. Suppose my friend Peter and I disagree about a claim in metaphysics—say, about whether abstract objects exist. While we are both professional philosophers, I am slightly his epistemic inferior on this question, since his primary research is in metaphysics and mine is in epistemology. He's studied the question more, read more papers on the topic, and has thought about the pros and cons of each side longer. I've taken a few metaphysics classes, and

thought about the question a bit, but not as much as he has. If the other conditions for peerhood are met (we genuinely disagree and are aware of the disagreement), and conciliationism is true, then Peter should conciliate—to an extent—after encountering our disagreement, but maybe not as much as he would with another philosopher who also specializes in metaphysics. Suppose instead that Peter disagrees with someone who has just been introduced to the debate about abstract objects. This person is even more of Peter’s epistemic inferior than I am, so Peter should conciliate even less. In a third case, suppose Peter disagrees with a senior philosopher who has been thinking about abstract objects for her entire career. Then, Peter should conciliate even more than in the standard peer case. In general, the point is that conciliationism does not depend on the existence of perfect epistemic peers. While it might be theoretically useful to focus on this idealized case, in real life, we should proportion our conciliating to the extent that the other is a peer.

This gives us an additional reason to prefer credal-conciliationism to belief-conciliationism. Beliefs are significantly too coarse-grained to capture all these differences. If they can’t even capture how I should respond to disagreement when I believe *p* and my exact epistemic peer withholds, they surely cannot capture cases of disagreement with inferiors, superiors, novices, and experts. We need an attitude that is more fine-grained and flexible—and this is exactly what a credence framework provides. Thus, this is yet another reason to prefer option (iv) to option (iii)—in response to disagreement, we should change our credal attitudes, rather than our belief-attitudes.

To take stock: if conciliationism is true, then at least two factors affect how we should change our credences in response to disagreement. The first is the updating rule we should adopt—and we saw in the previous section some reasons that this might not be as simple as splitting the difference. It may take various forms if, e.g. those who disagree both have credences on the same side of 0.5. The second is the extent to which the disagreement is among peers—whether the other is slightly inferior, slightly superior, or something else. It’s also worth noting, especially in real life, conciliating may require a balance between efficiency and accuracy. For example, it may be difficult to know exactly how our epistemic credentials compare, and thus, exactly how much each of us should conciliate when we disagree. But we often have a rough idea, and rationally responding to disagreement shouldn’t require perfection. Either way we strike this balance, conciliationism is more complex than simple difference splitting with perfect peers.

6. Objections

I close by addressing two objections to (1*) and (2*). First, one might wonder what should happen in cases of *iterated disagreement*—if I continually come into contact with smart people who strongly disagree with me on whether *p*. If I keep conciliating so my credence in *p* gets very, very low, should I ever give up my belief that *p*? It seems potentially problematic, if, say, my credence in *p* is 0.00001 but I nonetheless continue to believe *p*.

The answer to this question depends on the *normative* relationship between belief and credence. Note that dualism is merely a descriptive claim: that belief doesn’t reduce to credence, and credence doesn’t reduce to belief. Nonetheless, there are various views of the relationship between *rational* belief and *rational* credence that are consistent with dualism. For example, the Lockean thesis says that someone rationally believes *p* if and only if they have a rational credence in *p* above some threshold. (Foley 1993: ch. 4, Locke 2013, Shear & Fitelson 2018, Dorst 2019, and Lee & Silvia forthcoming defend the Lockean thesis.) Dualists can consistent either affirm or deny the Lockean thesis.

Let's suppose the Lockean thesis is true, and that the threshold is say, 0.7. You may start off believing p and having a high credence in p . As you conciliate with smart people who disagree with you, your credence gets lower and lower. Once your credence dips below 0.7, then you should give up your belief. Nonetheless, you don't give up your belief for *disagreement* reasons, but simply for coherence reasons—if you believe p , your credence in p just shouldn't be *that* low.

One might worry this kind of picture wouldn't satisfy a steadfast⁷. Insofar as we want to capture the intuitions behind steadfastness, this might motivate us to set the Lockean threshold quite low—maybe somewhere in the 0.5 range (or even lower). The lower the threshold, the more we can remain steadfast in light of disagreement. But if the steadfast still finds this unsatisfying, they may want to reject the Lockean thesis altogether.

Views that reject the Lockean thesis entail that belief and credence are relatively normatively independent. These views allow for the possibility that rational agents can believe p but have a very low credence in p . While this may seem counterintuitive, there are some compelling arguments against the Lockean thesis—see Friedman (2013), Buchak (2014), Smith (2016), Jackson (2019-c, 2020-a). This paper provides additional reasons to think that either (i) the Lockean threshold is quite low or (ii) there isn't a Lockean threshold at all. Especially for those who already dislike the Lockean thesis, this is an attractive way to capture both what is intuitive behind steadfast and conciliatory views. Buchak (forthcoming), for example, argues that two people could have opposite beliefs, i.e. one believes p and the other believes not- p , but both have the same credence in p . Here, I won't take a stand on whether there is a Lockean threshold, but I do think this paper provides reasons for skepticism. Alternatively, a more conciliatory version of my view combines (1*) and (2*) with a Lockean threshold. On this view, one can maintain one's beliefs in the face of disagreement—but only up to a certain point.

Second, one might wonder what my view says about cases of mundane disagreement, such as Christensen's restaurant case above. Should I really *continue to believe* everyone's share is \$43, even though you've calculated \$45? Is simply lowering my credence enough?

In response, there are two ways we can go. First, note that your credence that everyone's share is \$43 should not be extremely high, given your evidence is only that you've done the math in your head. Thus, maybe you can continue to believe in the face of a single case of disagreement, but if you get additional evidence you're wrong, you should probably give up your belief. Note that even if we accept (1*) and (2*) and reject the Lockean thesis, this doesn't mean we should never give up our beliefs due to disagreement, especially if (a) we didn't have strong evidence in favor of the proposition to begin with and (b) we have both disagreement and non-disagreement evidence against the proposition.

A second possibility is to say that (1*) and (2*) apply to beliefs that are important to us and central to our life commitments—beliefs that are at the core of our web of belief. However, in cases of mundane disagreement, like Christensen's restaurant case, conciliation with *both* belief and credence is appropriate. This is similar to the “weak conciliationism” proposed by Pittard (2019), on which significant conciliationism is called for in cases of ordinary disagreement, but one can remain steadfast in cases of “deep disagreement”—disagreement regarding one's central life commitments, e.g. religious, political, and moral views. My view captures that one should nonetheless conciliate *in some sense* in cases of deep disagreement, while continuing to believe the propositions central to their commitments and life projects. Buchak (forthcoming: Sec. 7) argues that many of the epistemic benefits of steadfastness don't apply in cases of mundane disagreement:

⁷ Thanks to Kevin McCain, Sofia Bokros, and Francesco Praolini.

“you don’t care that much about maintaining a correct belief on this topic over time, you don’t have a lot of other beliefs that depend on presupposing this calculation in your reasoning, you are not already involved in a long-term course of action on its basis, and so forth.” Thus, if we limit the scope in which (1*) and (2*) apply, this may enable us to both capture what’s intuitive about cases of mundane disagreement and also reap the epistemic benefits of steadfastness. Nonetheless, some steadfasters might worry that this picture is too conciliatory, and in that case, one can maintain that (1*) and (2*) apply across the board, to all cases of disagreement.

7. Conclusion

I’ve argued that disagreement leads to an epistemic dilemma—there’s good reason to think we should both change and not change our opinions when we encounter smart, well-informed people who disagree with us. I’ve argued that one way out of this dilemma is to maintain that we should change our credences, but not our beliefs, in response to disagreement. In addition to resolving this dilemma, I’ve provided two additional arguments for this view: (i) credal-conciliation allows for a flexibility that belief-conciliation doesn’t, creating space for more nuanced updating rules when simplistic ones won’t do, and (ii) credal-conciliation enables us to capture the epistemic significance of disagreeing with non-peers. I’ve also addressed two objections to my view. While I don’t take myself to have addressed every possible worry, I hope to have shown that this view is attractive and deserves a place at the table.

Acknowledgments: Thanks to Sofia Bokros, Kevin McCain, Francesco Praolini, and Scott Stapleford for helpful written comments on earlier drafts of this paper. Thanks to audiences at the Social (Distance) Epistemology Workshop, the University of Cologne/CONCEPT, the 2021 Eastern APA, and my “Disagreement and Permissivism” Foundations seminar at the Australian National University for helpful discussion.

Works Cited:

Ballantyne, Nathan. (2014). “Counterfactual Philosophers.” *Philosophy and Phenomenological Research* 88(2): 368–387.

Buchak, Lara. (2014). “Belief, Credence, and Norms.” *Philosophical Studies* 169(2): 285–311.

Buchak, Lara. (Forthcoming). “A Faithful Response to Disagreement.” *The Philosophical Review*.

Christensen, David. (2007). “The Epistemology of Disagreement: The Good News.” *The Philosophical Review* 116(2): 187–217.

Christensen, David. (2009). “Disagreement as Evidence: The Epistemology of Controversy.” *Philosophy Compass* 4/5: 756–757.

Christensen, David. (2016). “Conciliationism, Uniqueness, and Rationality Toxicity.” *Noûs* 50(3): 584–603.

Christensen David and Jennifer Lackey, eds. (2013). *The Epistemology of Disagreement: New Essays*. Oxford: OUP.

- Conee, Earl. (1994). "Against an Epistemic Dilemma." *Australasian Journal of Philosophy* 72(4): 475–481.
- Conee, Earl (2010). "Rational Disagreement Defended." In Richard Feldman and Ted Warfield (eds.), *Disagreement* (OUP): 69–90.
- de Finetti, B. (1974). *Theory of Probability*, vol. I. New York: Wiley.
- Dorst, K. (2019). Lockeans Maximize Expected Accuracy. *Mind* 128(509), 175–211.
- Eriksson, L. & A. Hájek. (2007). "What are Degrees of Belief?" *Studia Logica* 86(2), 183–213.
- Easwaran, Kenny, Luke Fenton-Glynn, Christopher Hitchcock, & Joel D. Velasco. (2016). "Updating on the Credences of Others: Disagreement, Agreement, and Synergy." *Philosophers' Imprint* 16(11): 1–39.
- Elga, Adam. (2007). "Reflection and Disagreement." *Noûs* 41(3): 478–502.
- Feldman, Richard. (2007). "Reasonable Religious Disagreements." In *Philosophers Without Gods: Meditations on Atheism and the Secular Life* (Louise Antony, ed.), pp. 194–214. Oxford: OUP.
- Fleisher, Will. (Forthcoming). "How to Endorse Conciliationism." *Synthese*.
- Foley, R. (1993). *Working without a Net*. Oxford: OUP.
- Friedman, J. (2013). Rational Agnosticism and Degrees of Belief. In: *Oxford Studies in Epistemology* 4. Ed. by Tamar Szabó Gendler and John Hawthorne, pp. 57–81. Oxford: OUP.
- Fritz, James. (2018). "Conciliationism and Moral Spinelessness." *Episteme* 15(1): 110–118.
- Jackson, Elizabeth. (2020-a). Belief, Credence, and Evidence. *Synthese* 197(11): 5073–5092.
- Jackson, Elizabeth. (2020-b). The Relationship Between Belief and Credence. *Philosophy Compass* 15(6): 1–13.
- Jackson, Elizabeth. (2019-a). "Belief and Credence: Why the Attitude-Type Matters." *Philosophical Studies* 176(9): 2477–2496.
- Jackson, Elizabeth. (2019-b). "How Belief-Credence Dualism Explains Away Pragmatic Encroachment." *The Philosophical Quarterly* 69(276), 511–533.
- Jackson, Elizabeth. (2019-c). Belief, Credence, and Faith. *Religious Studies* 55(2), 153–168.
- Jeffrey, R. (1965). *The Logic of Decision*. (second edition 1983). Chicago, IL: University of Chicago Press.
- Kelly, Thomas. (2005). "The Epistemic Significance of Disagreement." *Oxford Studies in Epistemology* (vol. 1), pp. 167–197. Tamar Gendler and John Hawthorne, eds. Oxford: OUP.

- Kelly, Thomas. (2008). "Disagreement, Dogmatism, and Belief Polarization." *Journal of Philosophy* 105(10): 611–633.
- Kelly, Thomas. (2010). "Peer Disagreement and Higher Order Evidence." In Richard Feldman and Ted Warfield (eds.), *Disagreement* (OUP): 111–174.
- Kelly, Thomas. (2013). "Disagreement and the Burdens of Judgment." In David Christensen and Jennifer Lackey (eds.), *The Epistemology of Disagreement: New Essays* (OUP): 31–53.
- King, Nathan. (2012). "A Good Peer is Hard to Find." *Philosophy and Phenomenological Research* 85(2): 249–272.
- Kitcher, Philip. (1990). "The Division of Cognitive Labor." *Philosophy of Science* 87: 5–22.
- Kitcher, Philip. (1993). *The Advancement of Science*. New York: OUP.
- Kornblith, Hilary. (2010). "Belief in the Face of Controversy." In R. Feldman and T. Warfield (eds.), *Disagreement*, pp. 29–52. Oxford: Oxford University Press.
- Lackey, Jennifer. (2008). "A Justificationist View of Disagreement's Epistemic Significance." In *Social Epistemology*, pp. 145–154. Alan Millar Adrian Haddock & Duncan Pritchard, eds. Oxford: OUP.
- Lee, Matthew & Paul Silva Jr. (Forthcoming). "Toward a Lockean Unification of Formal and Traditional Epistemology." *Episteme*. doi:10.1017/epi.2020.11.
- Locke, D. (2014). The Decision-Theoretic Lockean Thesis. *Inquiry* 57, 28–54.
- Lougheed, Kirk. (2020). *The Epistemic Benefits of Disagreement*. New York: Springer.
- Moon, A. (2019). A New Puzzle About Belief and Credence. *The Canadian Journal of Philosophy* 49(2), 272–291.
- Muldoon, Ryan. (2013). "Diversity and the Division of Cognitive Labor." *Philosophy Compass* 8(2): 117–125.
- Pettit, Philip. (2006). "When to Defer to Majority Testimony—And When Not." *Analysis* 66(3): 179–87.
- Pittard, John. (2019). *Disagreement, Deference, and Religious Commitment*. Oxford: OUP.
- Ramsey, F. (1926). "Truth and Probability." In: *The Foundations of Mathematics and Other Logical Essays*. Ed. by R. B. Braithwaite, Ch. VII, pp. 156–198. London: Kegan, Paul, Trench, Trubner & Co.
- Rosen, Gideon. (2001). "Nominalism, Naturalism, Epistemic Relativism." *Philosophical Perspectives* 15 (*Metaphysics*): 69–91.
- Ross, J., and M. Schroeder. (2014). "Belief, Credence, and Pragmatic Encroachment." *Philosophy and Phenomenological Research* 88(2): 259–288.

- Russell, J. S., Hawthorne, J., and Buchak, L. (2015). Groupthink. *Philosophical Studies*, 172(5): 1287–1309.
- Sampson, Eric. (2019). “The Self-Undermining Arguments from Disagreement.” *Oxford Studies in Metaethics* 14: 23-46.
- Shear, T. and B. Fitelson. (2019). Two Approaches to Belief Revision. *Erkenntnis* 84(3), 487–518.
- Sherman, Benjamin. (2015). “Questionable Peers and Spinelessness.” *Canadian Journal of Philosophy* 45(4): 425-444.
- Schupbach, Jonah. (2018). “Troubles for Bayesian Formal Epistemology? A Response to Horgan.” *Res Philosophica* 95(1): 189–197.
- Schwitgebel, E. (2019). Belief. In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. <http://plato.stanford.edu/entries/belief/>
- Smith, M. (2016). *Between Probability and Certainty: What Justifies Belief*. Oxford: OUP.
- Staffel, J. (2017). “Accuracy for Believers.” *Episteme* 14(1): 39-48.
- Turnbull, Margaret & Sampson, Eric. (2020). "How Rational Level-Splitting Beliefs Can Help You Respond to Moral Disagreement." In *Higher-Order Evidence and Moral Epistemology* (M. Klenk ed.). New York: Routledge.
- van Inwagen, Peter. (1996). “Is it Wrong Everywhere, Always, and for Anyone to Believe Anything on Insufficient Evidence?” In *Faith, Freedom and Rationality* (Jeff Jordan and Daniel Howard-Snyder, eds.), pp. 137-154. Savage, Maryland: Rowman and Littlefield.
- Warfield, Ted & Richard Feldman. (2010). *Disagreement*. Oxford: OUP.
- Weisberg, J. (2020). “Belief in Psyontology.” *Philosopher’s Imprint* 20(11): 1–27.