

This paper appears in *Synthese* 198: 4985–5008. (This is a special issue on the work of Catherine Z. Elgin.) The published version can be found here: <https://link.springer.com/article/10.1007/s11229-019-02384-6>.)

ELGIN'S COMMUNITY-ORIENTED STEADFASTNESS

ABSTRACT: In recent years, epistemologists have devoted enormous attention to this question: what should happen when two epistemic peers disagree about the truth-value of some proposition? Some have argued that that in all such cases, both parties are rationally required to revise their position in some way. Others have maintained that, in at least some cases, neither party is rationally required to revise her position. In this paper, I examine a provocative and under-appreciated argument for the latter view due to Catherine Z. Elgin (2010, 2012, 2017, and 2018). I defend it against a series of objections, and I then identify some fruitful ways in which her view could be developed further.

Catherine Z. Elgin holds that the existence of certain kinds of peer disagreement can help communities of inquirers achieve their shared epistemic goals. Accordingly, she thinks that it can be rational to stand one's ground in such disagreements. In Section 1, I set out two main competing views about what rationality licenses in these cases. In Section 2, I explain why and how Elgin wishes to reframe this debate. In Section 3, I introduce her argument for her view, and in Section 4, I offer some support for it. In Section 5, I defend Elgin's view against seven objections. Finally, in Section 6, I identify eight fruitful ways in which her view could be developed further.

1. FRAMING THE DEBATE

In recent years, epistemologists have devoted enormous attention to disagreement between acknowledged epistemic peers.¹ Here is one definition of peerhood:

Let us say that two individuals are epistemic peers with respect to some question if and only if they satisfy the following two conditions: (i) they are equals with respect to their familiarity with the evidence and arguments which bear on the question, and (ii) they are equals with respect to general epistemic virtues such as intelligence, thoughtfulness, and freedom from bias (Kelly 2005, 174–5).

Drawing on definitions like this, Nathan King (2012) observes that the contemporary literature about the epistemology of disagreement has focused primarily on cases where the following four conditions are satisfied:

- (a) The disagreement condition: S believes P, while T believes \sim P;
- (b) The same evidence condition: S and T have the same P-relevant evidence, E;
- (c) The dispositional condition: S and T are equally disposed to respond to E in an epistemically appropriate way; and
- (d) The acknowledgement condition: S and T have good reason to think conditions (a)-(c) are satisfied.

So the discussion has primarily sought to determine what rationality licenses, and what it prohibits, when two acknowledged peers learn that they disagree about the truth-value of some proposition. Some say that in all such cases, both S and T possess a defeater for their belief, and so neither can rationally maintain their original confidence in it.² Others say that at least in some such cases, S and T can be rational in sticking to their guns concerning P.³

But King persuasively shows that this construal of the debate is unduly narrow. In the real world, two inquirers rarely possess the *exact same* evidence, nor are they *equally disposed* to respond to it appropriately, nor does either party generally have *good reason* to think the other has the exact same evidence and equally appropriate dispositions. So conditions (b), (c), and (d) are rarely satisfied. Accordingly, King recommends that epistemologists worry less about symmetry, and focus more on what is *unclear*, in cases of disagreement. In particular, King proposes that the debate be broadened to focus on principles like this instead:

- (SK1) If S believes P and is aware of some other subject T who believes \sim P, and S has no reason to think S's total epistemic position with respect to P renders S more likely to be correct than T, then S is not [epistemically] rational in believing P [as confidently as S did prior to becoming aware of T: S should reduce her confidence in P].⁴

King rightly claims that the antecedent of (SK1) is more frequently satisfied than conditions (b)-(d) jointly are, and that, accordingly, focusing on (SK1) will appropriately broaden the debate. Suppose that (SK1) is true, and that S recognizes that it is true: rationality then requires S to reduce her credence in P to some degree, perhaps even to the point of suspending judgment about P altogether. For now, let's say that that *conciliationism* affirms exactly that, while the *steadfast view* rejects (SK1).

2. REFRAMING THE DEBATE

Catherine Elgin identifies drawbacks of both views. According to her, conciliationism leads to epistemic spinelessness, for the moment that S discovers that her overall epistemic position renders her no more likely to be correct about P than T, rationality obliges her to reduce her credence in P (2010, 54; 2018, 14-15). Meanwhile, Elgin says, the steadfast view can make S stubborn by making her insufficiently sensitive to counter-evidence. That's because S's first-order reasons for believing P constitute reasons for thinking that evidence *against* P is *misleading*, and given this, S is licensed to disregard it. Such stubbornness, Elgin

says, can deprive S of valuable epistemic resources, which makes it epistemically irresponsible (2017, 6). For example, to the extent that S is stubborn, she might no longer use her engagement with T's view to identify and correct mistakes, to re-examine her own positions, to develop better methods of assessing the evidence, etc. (2010, 55). The fundamental problem with both positions, says Elgin, is that they discourage S from really, seriously engaging with T's view (2018, 14-15).

Accordingly, Elgin devises a new account of what S should do in such cases. She begins by urging that this debate's focus on *belief* is misguided, since beliefs are involuntary:

I may find myself with a belief suspended as a result of evidence, argument, testimony, or disagreement; or I may find my belief unmoved by evidence, argument, testimony, or disagreement. But my response is not under my control. Debates about whether I should suspend belief in the face of peer disagreement are wrong-headed. They are like debates about whether I should be less than six feet tall. I don't have any choice (2010, 61).

Invoking a distinction due to L. Jonathan Cohen (1992), Elgin urges that the debate about peer disagreement should concern *acceptance*, rather than belief, since it is a voluntary doxastic state. Elgin says that "To accept that p is to adopt a policy of being willing to treat p as a premise in assertoric inference or as a basis for action where our interests are cognitive" (2010, 64).⁵ Although Elgin does not say so, it's important to note that this sort of willingness comes in degrees: one can be more or less willing to φ . Accordingly, acceptance is also a degree notion. Given all this, on Elgin's view the debate between steadfasters and conciliationists should really be about:

(SK1*) If S accepts P and is aware of some other subject T who accepts $\sim P$, and S has no reason to think S's total epistemic position with respect to P renders S more likely to be correct than T, then S is not epistemically rational in accepting P to the degree S did prior to becoming aware of T: S should reduce her degree of acceptance in P.

Hereafter, then, I take *conciliationism* to be expressed by (SK1*), and the *steadfast view* to be its denial. Elgin thinks that once the debate is reframed in this way, an important reason for resisting (SK1*) can emerge – one that avoids the very drawback of steadfastness that she identifies.

3. ELGIN'S COMMUNITY-ORIENTED STEADFASTNESS

Elgin thinks that when the antecedent of (SK1*) is satisfied, to refrain from accepting P is, *ceteris paribus*, a "cognitively impoverishing stance" (2010, 65). To illustrate this idea, she offers the story of Jack, a paleontologist who has become convinced of conciliationism. Jack considers a wide range of propositions whose truth-value is contested within the relevant community of paleontologists. Since he is a conciliationist, his reflection on these disagreements leads him to

refrain from accepting a large number of propositions that he had hitherto accepted. Jack now has substantially fewer premises available for use in his reasoning concerning paleontology. And if all experts in this domain were likewise to refrain from accepting such controversial claims, the result would be catastrophic for inquiry: it would leave “the paleontological community with few premises about their subject matter, yielding a sparse and moth-eaten fabric of cognitive commitments” (2010, 65, and see also 2012, 80-1).

The stage is thus set for Elgin’s argument for steadfastness, which she expresses as follows:

When the reasons favoring each side of a dispute are sparse or exceedingly delicate, or the evidence is equivocal, or each side can solve important common problems that the other cannot, it may be better for the epistemic community that both positions continue to be accepted ... Each group then can draw on a different range of commitments for premises in their reasoning and as a basis for their actions. By developing their positions, they put them to the test ... Tolerance of disagreement is epistemically valuable only when the disagreement is among parties who have sufficient expertise in an area that their opinions are individually worth accepting, and where the evidence at hand is equivocal. When these conditions are met, a community of inquiry may best be served if epistemic peers resolutely reason and act on opinions about which other equally competent inquirers disagree. In such cases, peers who disagree have reason to consider each other wrong but not irrational. Perhaps in the fullness of time the disagreement will be resolved. That remains to be seen (2010, 67-68).⁶

Elgin’s central idea is that in certain cases, a community’s shared epistemic goals may be better served if individual inquirers persist in accepting their views in the face of the relevant kinds of disagreement. Her view, then, can properly be termed *community-oriented steadfastness*.

4. BENEFITS OF VIEWPOINT DIVERSITY FOR COMMUNITIES

Elgin thinks that that steadfastness can promote the existence of a wider range of acceptances within the community, and that in certain cases, this viewpoint diversity can help a community of inquiry to achieve its epistemic goals. This idea is, in fact, rather well-supported by social-scientific research. It also comports rather well with analyses of the division of cognitive labour in the philosophy of science, and with a certain evolutionary account of the development of reasoning. I’ll briefly discuss each in turn.

Although Elgin doesn’t mention it, her claim that viewpoint diversity can benefit communities of inquiry enjoys considerable empirical support.⁷ For one thing, viewpoint diversity mitigates the intellectually pernicious effects of a well-established phenomenon in collective inquiry: group polarization (Moscovici and Zavalloni 1969; Lamm and Myers 1978; Vinokur and Burnstein 1978; and Sunstein 2002). Furthermore, decades of psychological research suggest that when groups contain diverse viewpoints, they perform better in a wide range of creative,

problem-solving, predictive, and discovery-based tasks than do individuals and homogenous groups. (See, for example, Moscovini and Personnaz 1980; Hastie 1986; Laughlin and Ellis 1986; Nemeth 1986; Nemeth 1995; Gigone and Hastie 1997; Kuhn et al. 1997; Mosheman and Geil 1998; Jehn et al. 1999; Brodbeck et al. 2002; Greitemeyer et al. 2006; Schulz-Hardt et al. 2006; van Knippenberg and Schippers 2007; Crisp and Turner 2011; Krause et al. 2011; Trouche et al. 2014; Duarte et al. 2015; and Woolley et al. 2015.)⁸

Since groups with robust viewpoint diversity typically perform better in a wide range of intellectual tasks than do individuals or homogenous groups, it would not be surprising if scientific inquiry attempted to harness this fact. And indeed, this is just what philosophers of science have postulated. The basic idea here is that just as research teams working on a specific problem can benefit from viewpoint diversity within the team, so too a larger community of inquiry can benefit from the division of cognitive labour by having many research teams work on the same problem from different starting points and in varied ways (while, of course, sharing their results and learning from each other).⁹

Moreover, Elgin's view fits rather nicely with the *argumentative theory of reasoning*, according to which the primary evolutionary function of reasoning is to devise and evaluate arguments that can persuade others (Sperber 2001; Mercier and Sperber 2011; 2017, 211-221).¹⁰ One surprising – and illuminating – aspect of this theory involves confirmation bias. This bias is generally thought to be epistemically undesirable, but, according to the argumentative theory of reasoning, confirmation bias can in fact promote the division of cognitive labour (which, as we have seen, is thought to benefit inquiry). Mercier and Sperber explain: "When a group has to solve a problem, it is much more efficient if each individual looks mostly for arguments supporting a given solution. They can then present these arguments to the group, to be tested by the other members ... This joint dialogic approach is much more efficient than one where each individual on his or her own has to examine all possible solutions carefully" (2011, 65). Surveying recent empirical work on this phenomena, Hallsson and Kappel (forthcoming) find that this is exactly what the studies show:

Motivated reasoning [the tendency to construe evidence in order to support a conclusion that one already believes or wants to believe] can increase one's ability to find good reasons in favor of one's view, and to critique reasons against it, compared to more dispassionate reasoning. When both sides of an issue are represented in deliberation, this results in a broad range of reasons being considered for each side, and the selective retention of the best ones. As a result, members of deliberating groups composed of people who disagree tend to arrive at beliefs that are better supported by the evidence than individuals reasoning alone, or members of groups composed of people who agree prior to deliberation (2).

All of this supports Elgin's claim that steadfastness, at least in certain cases, can indeed benefit communities of inquiry.¹¹ But, returning to the debate at hand, does this mean that Elgin's argument against (SK1*) succeeds? To tackle this question, I now consider seven objections to Elgin's position.

5. OBJECTIONS TO ELGIN'S VIEW

In this section, I discuss seven objections that might be offered against Elgin's rejection of (SK1*), and I suggest that none of them is decisive. I consider them in what I take to be increasing order of strength.

5.1. *Elgin Begs the Question*

One might worry that Elgin begs the central question at issue when she explicitly stipulates that the disagreeing parties in the relevant cases each have opinions that are "individually worth accepting" (2010, 68); that they are "epistemically acceptable alternatives" (2018, 18). The conclusion of her argument is that in certain conditions, at least one party to the disagreement can be epistemically rational in maintaining her degree of acceptance, so it will not do to stipulate this at the outset. This charge, however, is easily avoided. I take it that she wants to restrict the extent to which her view licenses steadfastness. In particular, I take it that she does not want to license the maintaining of absurd views in the face of relevant disagreement. She is clear that "not all points of view merit attention" (2018, 18); and that her position is not one of "mindless tolerance of any opinion whatsoever" (2012, 82, and see also 2017, 110, 114). To this end, she could stipulate that opinions to which her steadfastness applies must be 'plausible', or 'prima facie worth accepting', or even 'worth accepting prior to considering disagreement'. Some such restriction will ensure that steadfastness concerning absurd views is not licensed by her view, while avoiding begging the question in her favour.¹²

5.2. *Elgin's View Licenses Stubbornness*

As we have seen, Elgin thinks that rejecting the original (SK1) can license objectionable stubbornness. That's because in so doing, S can use her first-order reasons for believing P to dismiss T's evidence against P as misleading. And to do this, Elgin says, is to shirk one's epistemic responsibility: "Rationality requires that epistemic agents be open-minded; they should not intentionally blind themselves to evidence" (2018, 14). Someone might say, however, that the very same problem infects Elgin's rejection of (SK1*). After all, it seems that S could just as well use her first-order reasons for *accepting* P to dismiss T's reasons for not accepting P as misleading, which would be just as stubborn. But Elgin has the resources to reply to this charge. She can say – and indeed does say – that continuing to accept P in the face of disagreement *for the sake of one's community* does not in fact license this sort of stubbornness. On the contrary: "It requires [S] to be open to, and perhaps even to seek out, emerging counter-evidence and counter-arguments ... It undermines complacency and fosters intellectual respect" (2018, 20, and see also 2012, 82). All of this is precisely what's involved in S's continuing to accept P for the sake of her epistemic community, according to Elgin.¹³

5.3. *Diverse Acceptances Aren't Needed for Community-Oriented Benefits*

Aristotle is rumoured to have said that it is a mark of an educated mind to be able to *entertain* a thought without *accepting* it.¹⁴ Now, whether or not Aristotle actually said anything of the sort, the distinction is worth considering in this

context. That's because one might object that the community-oriented epistemic benefits of viewpoint diversity can be achieved without licensing inquirers to remain steadfast in their *acceptances* in situations described by the antecedent of (SK1*). In particular, one might think that epistemically fruitful diversity could be achieved by having members of the community merely *entertain P* in such situations – while nevertheless conciliating – rather than remaining steadfast by continuing to accept P with undiminished confidence.¹⁵ Recall that for Elgin, to accept P is “to adopt a policy of being willing to treat p as a premise in assertoric inference or as a basis for action where our interests are cognitive” (2010, 64). In certain contexts, at least, it might be possible to entertain P (perhaps as a provisional hypothesis) without being willing to treat it as a premise or as a basis for action in the manner that Elgin suggests.

Interestingly, empirical research suggests that stronger attitudes than entertaining are needed to bring about the desired benefits to group inquiry. For example, Nemeth et al. (2001) compared the effect of genuine dissent within groups to the effect of the presence of a ‘devil’s advocate’ (DA) in relevantly dissent-free groups – and they concluded that groups with *genuine* dissent perform better. Nemeth et al. speculate that: “[a]rmed with the belief that they have considered alternatives by virtue of exposure to the DA, people may become even more convinced of the truth of their initial position – and possibly more rigid and resistant to reconsideration” (709). Relatedly, Schultz-Hardt et al. (2002) concluded that in group inquiry, genuine “heterogeneity was more effective in preventing a confirmatory information-seeking bias than devil’s advocacy was” (563).¹⁶ These results suggest that a more committed epistemic stance, like acceptance, is required in order to realize the relevant sorts of community-oriented benefits.

5.4. Elgin’s View Subordinates Individual Epistemic Interests to Group Needs

One might object that Elgin’s steadfastness inappropriately subordinates individual epistemic interests to group needs. It is easy to imagine an inquirer saying something like this: “Sure, I do care about whether my community ultimately figures out whether P, but, look, I care rather more about *my* getting to the truth about P.” Moreover, it is easy to imagine cases in which this is an eminently reasonable position, given one’s resources and cognitive interests. So, isn’t Elgin demanding too much with her focus on the potential benefits for epistemic communities?

Well, no. First, it may be that an epistemic community’s goals are simply constituted by shared goals of its individual members – and if so, there simply can be no tension between the two. Second, even if a community’s goals are not simply constituted by the shared goals of individuals, it may be that the best way for S to achieve her epistemic goal of getting to the truth about P just is by remaining steadfast, thereby promoting her community’s intermediate goal of achieving viewpoint diversity, which in turn promotes the community’s ultimate goal of figuring out whether P.¹⁷

The deeper problem with this objection is that it claims that Elgin’s view is too demanding for individual inquirers – but Elgin’s view doesn’t actually make any demands at all. In particular, Elgin does not say that one rationally *ought* to

remain steadfast in the relevant kinds of cases; she merely argues that one is not rationally required to revise. So Elgin's view cannot be read as inappropriately subordinating the epistemic needs of the individual to those of her community.

5.5. *Elgin's View Cannot Defeat Conciliationism Based on Possible Disagreement*

The dispute between conciliationists and steadfasters generally involves cases in which both parties to the disagreement are supposed to be *actual*. But some conciliationists have argued that reflection on disagreement with merely *possible* peers should likewise lead to a reduction in confidence in one's views.¹⁸ The basic idea is this: conciliationism is fundamentally motivated by the thought that there might be good reasons for thinking otherwise than you do, and it doesn't really matter whether these reasons are actually expressed by anyone.¹⁹ Someone persuaded by this might be tempted to think that Elgin's steadfastness fails to count against it.²⁰ Such a person might point out that Elgin's defence of steadfastness appeals to the epistemic interests of *actual* communities of inquirers, and then insist that these considerations are inert against conciliationist appeals motivated by merely *possible* disagreement.

This objection fails, however, since Elgin can plausibly respond in the following way.²¹ S may initially be tempted by the thought that she should reduce her confidence in P after realizing that her overall epistemic position is not superior to that of some possible agent, T, who accepts $\sim P$. But S may well go on to reason like this:

Even though, to the best of my knowledge, T is not an actual person, my actual community of inquirers may well benefit from my remaining steadfast in accepting P in the face of my disagreement with T. After all, my willingness to employ P in the relevant inferences and actions may help my community figure out once and for all whether P. This could happen in several ways. Here are two examples. First, my continued acceptance of P might ultimately help *me* to find the fatal flaw in T's reasons for accepting $\sim P$, and this could well constitute, or lead to, genuine epistemic progress for my community as a whole. Second, my acceptance of P might spur *others* in my actual community of inquirers to attend to this issue more carefully, and this might help lead them, and eventually all of us, to finally know the truth-value of P.

Of course, things could go precisely the other way: S's acceptance of P might lead her actual community of inquirers to go collectively in exactly the wrong direction with respect to the truth-value of P. But, for the moment, this is irrelevant.²² If there is at least a live possibility that S's continued acceptance of P in the face of possible disagreement will confer adequate epistemic benefits on S's community of inquirers, then Elgin can appeal to it, exactly as she does in the case of disagreements with actual peers.

5.6. Appealing to Community-Oriented Reasons is not Epistemically Rational

Christensen (2007) pre-dates Elgin's arguments, but he anticipates one motivation for her view when he writes: "It's quite plausible that knowledge is best advanced by people exploring, and attempting to defend, a variety of answers to a given question. Perhaps, human psychology makes this easier to do when investigators actually have a lot of confidence in the hypotheses they're trying to defend" (215). Christensen, however, also raises a worry: "Certain sorts of inquiry might well work best when a variety of investigators have irrationally high levels of confidence in a variety of pet hypotheses. So there may well be important epistemic benefits to certain patterns of irrational [acceptance]. But I would argue that the patterns of [acceptance] are no more *epistemically* rational for all that" (215, emphasis added).²³ If this were offered as an objection to Elgin's view, it would fail. The problem, of course, is that in this dialectical context it is question-begging to merely *stipulate* that the relevant acceptances are not epistemically rational. The central aim of Elgin's argument, after all, is to *show* that in some cases, it is epistemically rational to maintain one's acceptances in the face of disagreement.

There are, however, more plausible ways to press this objection. In the remainder of this section, I distinguish three such ways, and respond to each in turn.²⁴ The first way holds that there is something epistemically illicit about sticking to one's guns in the face of peer disagreement for the sake of community-oriented benefits down the road. Such considerations are *pragmatic* or *prudential*, one might say, but they are not properly epistemic.²⁵

In response, it is worth noting that there are now quite a few attempts to develop accounts of how steadfastness for community-oriented reasons can be epistemically appropriate.²⁶ While the details differ, these accounts all make a common point: community-oriented reasons for steadfastness are indeed *epistemic*, in that they seek to bring about benefits to communities of inquiry that concern epistemic desiderata like justification, truth, belief, and knowledge. Jon Matheson, for example, distinguishes *synchronic* justification from *diachronic* justification, and suggests that community-oriented steadfastness invokes the latter, and that it is a properly epistemic notion (2015a; 2015b, 146-147). Relatedly, Kirk Lougheed and Robert Simpson argue that community-oriented reasons are *indirect epistemic reasons* (2017).²⁷ Space does not permit me to evaluate all these proposals, but I mention them here in order to show that it will not do to simply dismiss community-oriented reasons as non-epistemic. Serious accounts of such reasons as epistemic have been given (in addition to Elgin's), and critics would do well to engage with them.

Here is a second way to press this line of objection. Even if community-oriented reasons are epistemic, it might seem that defending steadfastness by appeal to them objectionably subsumes epistemic rationality under instrumental rationality. In order to clarify this charge, I will briefly set out one of Thomas Kelly's (2003) important criticisms of *instrumentalism*: the view that epistemic rationality is really just a species or special case of instrumental rationality. Kelly considers the following vivid case: S actively desires *not* to know how a given movie ends, since this would spoil her enjoyment of the movie at some later date. S implores her friend T not to give away the ending. T, however, blurts it out. Given that S knows that T is a reliable testifier on such matters, S immediately learns how

the movie ends, thereby gaining the very knowledge that she wished to avoid. Kelly thinks it is epistemically rational for S to believe (or, we can add, accept) the relevant proposition about the ending, but her doing so failed to serve any of her goals; in fact, it thwarted her stated goal. Kelly intends this to be a counterexample to the thesis that all epistemic rationality is instrumental rationality. And it might seem that if Kelly is right, this poses a problem for Elgin's view, insofar as it assumes instrumentalism.

Is Kelly right? Adam Leite complains that Kelly's construal of the relevant goal is too narrow, and, in particular, too short-term. He suggests that the instrumentalist can say that "epistemic rationality is a matter of forming beliefs in accordance with a system of habits or procedures whose consistent application would bring about the optimal satisfaction of one's cognitive goals over time" (2007, 461). It's plausible to think that being sensitive to the evidence of reliable testifiers is indeed an important part of such a system, even though such sensitivity can sometimes thwart one's short-term goals. Leite intends this move to open up space for the instrumentalist to claim that epistemic rationality is, after all, a species or special case of instrumental rationality.²⁸ Leite is quite right to say that, in principle, the instrumentalist can avoid objections like Kelly's by appealing to more general, longer-term goals (and associated, habits, procedures, and systems). After all, the instrumentalist needn't say that epistemic rationality serves one's *every* epistemic goal, nor even that specific short-term epistemic goals are more important than more general, long-term goals. Elgin might even add that some of these latter goals are community-oriented; this would fit well with the spirit of her project. In short, then, a broader account of the relevant goals can help Leite – and Elgin – to defang this objection to instrumentalism levelled by Kelly.²⁹ That said, a complete vindication of instrumentalism along these lines will, of course, require much more to be said about these goals, habits, procedures, and systems.

For my purposes, however, it's important to see that whatever the (de)merits of instrumentalism, Elgin's defence of steadfastness doesn't presuppose this view at all. Elgin simply aims to show that it can be epistemically rational, in certain situations, to remain steadfast in your acceptances for the sake of community-oriented epistemic benefits. This is compatible with the view that all epistemic rationality is instrumental, but does not require it. Perhaps some epistemic rationality is instrumental, and some is not. Accordingly, Kelly's arguments against instrumentalism, even if sound, pose no threat to Elgin's view.

Here is a third way to press the objection that community-oriented reasons cannot be epistemically rational. This way insists that (a) epistemic rationality is exclusively a matter of rational obligations, and, in particular, that (b) epistemic rationality includes an exceptionless obligation to base one's acceptances solely on the evidence available to one at that time. Given these claims, of course, it simply cannot be epistemically rational to stick to one's guns in the face of the relevant sort of disagreement for the sake of future community-oriented epistemic benefits.

Here are three points that Elgin can make in reply to this way of pressing the objection. The first is general: both (a) and (b) are enormously controversial positions in their own right, and as such, simply appealing to them cannot decisively refute her view.³⁰ The second point focusses on the debate about disagreement in particular. It is this. The denial of (a) and (b) is not unique to Elgin: virtually every defence of steadfastness in the disagreement literature claims

that in the circumstances described in the antecedents of (SK1) and (SK1*), it can be epistemically rational (i.e. epistemically permissible) for agents to stick to their guns.³¹ Indeed, defending the epistemic permissibility of so doing is the essence of the steadfast's project.³² So appealing to (a) and (b) does not specifically target Elgin's view.

Here is a final point worth mentioning. It's true that Elgin's view has a broadly consequentialist flavour: after all, she defends the epistemic permissibility of remaining steadfast in the relevant circumstances by appeal to future epistemic benefits for communities. Now, epistemic consequentialism typically rejects (a) and (b), and it is a widespread view.³³ Indeed, as Dunn and Ahlstrom-Vij note, "much of contemporary epistemology has been marked by ... a tacit endorsement of *epistemic consequentialism*" (2018, 4). Accordingly, Elgin's view comports well with a mainstream view.

I grant that these three points will not convince the diehard defender of (a) and (b). But in my view, in this dialectical context, these three points do jointly shift the burden of proof back to the defender of (a) and (b) to show that there is something illicit in Elgin's position. This does not mean, of course, that Elgin's defence of steadfastness on epistemic consequentialist grounds wholly succeeds. In fact, in Section 6 I set out eight desiderata for further work that will be required to flesh out her view. First, however, I turn to one final objection.

5.7. No One Can be Sure Enough about the Likelihood of Epistemic Benefits

The final objection to Elgin's view that I consider here is, I believe, the strongest. To see the motivation for it, we can begin by asking: what conditions must be satisfied in order for an agent to be entitled to remain steadfast, in the face of relevant disagreement, for community-oriented reasons? So far, we have seen two:

- (i) Elgin focuses on cases in which "... reasons favoring each side of a dispute are sparse or exceedingly delicate, or the evidence is equivocal, or each side can solve important common problems that the other cannot" (2010, 67), so it seems that S's position should satisfy at least one of these disjuncts.
- (ii) I suggested earlier that Elgin should also restrict her focus to disagreements in which each side's position surpasses some minimal epistemic threshold, in order to prevent her view from licensing steadfastness for absurd views.³⁴

But perhaps this is not enough. One might argue that a further condition must be met in order for S to be entitled to remain steadfast, namely:

- (iii) S reasonably accepts that her remaining steadfast concerning P for time interval t is, all things considered, *likely* to help her community achieve its relevant epistemic goals.

The argument for this condition's being necessary is fairly straightforward. If S has no acceptances whatsoever about the likelihood of her steadfastness (concerning P for time interval t) benefitting her community, then it seems that S cannot invoke

community-oriented reasons for her steadfastness. So she should have some acceptances about this matter. Moreover, it seems plausible to focus only on *reasonable* acceptances, since the epistemic goal here is to rationally license S's steadfastness. Given this, what options are there for S to reasonably accept about this likelihood? It seems that there are only four:

- S might reasonably accept that her remaining steadfast (concerning P for time interval t) is, all things considered, *unlikely* to help her community achieve its relevant epistemic goals.
- S might reasonably accept that her remaining steadfast (concerning P for time interval t) is, all things considered, *neither likely nor unlikely* to help her community achieve its relevant epistemic goals.
- S might reasonably be *agnostic* about the whole issue; and, finally,
- S might reasonably accept that her remaining steadfast (concerning P for time interval t) is, all things considered, *likely* to help her community achieve its relevant epistemic goals.

It is very tempting to say that only the last one of these can license S to remain steadfast for community-oriented reasons. If that's correct, the case for (iii) is complete.

But now comes the objection. Setting aside toy cases beloved by philosophers, it is exceedingly unlikely that any real agent, S, in any real community of inquirers, can ever reasonably accept that her remaining steadfast (concerning P for time interval t) is, all things considered, *likely* to help her community achieve its epistemic goals. After all, for S to be reasonable in accepting that such a thing is all-things-considered-likely, S would require reliable access to a vast store of information about the (likely) short-term, medium-term, and perhaps even long-term consequences of S's steadfastness during interval t , and how these (are likely to) bear on the community's overall pursuit of the truth concerning P. These would include, of course, consequences for S herself, consequences for each of her interlocutors (including T), consequences for her interlocutors' other interlocutors apart from S and T, and, further, consequences for those who are uninvolved in the dispute during t , but who are (directly or indirectly) affected by those who *are* involved during t in ways that ultimately bear on the community's goals. Et cetera. The number of causal chains concerning which S would need to have reasonable acceptances is staggering. S might not have to be omniscient to pull it off, but S would surely require extraordinary epistemic powers to even come close – powers that earthly agents do not possess.

If all this is plausible, the result is an unhappy one for Elgin's community-oriented steadfastness. It may indeed be that, in some case or other, S's community is likely to benefit from her remaining steadfast concerning P for time interval t . But if S can *never* reasonably accept as much, and if condition (iii) must be met for S to be rationally permitted to remain steadfast, then S is never entitled to appeal to community-oriented reasons in order to remain steadfast. And, accordingly, community-oriented reasons cannot support the steadfast view.³⁵

How should Elgin reply? I suspect that her best defence would be to argue that while condition (iii) is surely *sufficient* to license S's steadfastness,³⁶ it is not *necessary*, precisely because it places wholly unreasonable demands upon individual epistemic agents. It's just too demanding to expect S to *consider all things*, or even to consider all things relevant to whether her acceptance of P for time interval t is likely to help her community, and to then make reasonable judgments about what will happen in the future.

The objector could regroup, however, by proposing a weaker account of what is required for S to be reasonable in believing that her steadfastness will help her community, and arguing that this condition is rarely if ever satisfied. Here is such a condition:

- (iv) S reasonably accepts that her remaining steadfast concerning P for time interval t is, *given the relevant evidence she can reasonably be expected to acquire*, likely to help her community achieve its relevant epistemic goals.³⁷

The purpose of the italicized clause in (iv), of course, is to defang the worries just raised, by reducing the epistemic burden on S.

But isn't this burden still too heavy for S to shoulder? While S is no longer required to form reasonable judgements about what is likely to help her community, *considering all things*, she is still required to form reasonable judgments about what is in fact likely to help her community achieve its relevant epistemic goals, *given her available evidence*, and perhaps this is still too much to ask. The problem is that this still requires S to hold before her mind all her available evidence, to then to peer into some array of possible futures, and to make incredibly complex judgments about what will (or likely will) happen as a result of her steadfastness – and perhaps S just can't be expected to do that. If so, Elgin can reply that while (iv), like (iii), may well be a plausible sufficient condition for S to remain steadfast, it is not a necessary one. Again, it's just too demanding.

Here is yet another way to press the objection. If the fundamental problem with (iii) and (iv) is their focus on the future, maybe it's better to orient S towards the *past* instead. To motivate this idea, consider the striking historical case of the Hungarian physician Ignaz Semmelweis (1818-1865), which is discussed by Lougheed (2018a). Against the prevailing medical wisdom of his time, Semmelweis instructed his medical students to wash their hands before delivering babies. When they did, the incidence of puerperal (childbed) fever among the women they attended declined dramatically. Semmelweis, of course, knew nothing of germs, but his discovery of a correlation between hand-washing and the prevention of communicable disease – and, crucially, his steadfastness in the face of serious and sustained disagreement – ultimately led to significant epistemic benefits for the medical community and indeed for humanity generally.³⁸ Reflection on this sort of case might prompt the thought that (iv) should be replaced with:

- (v) S reasonably accepts that her remaining steadfast concerning P for time interval t is, given the relevant evidence she can reasonably be expected to acquire, *relevantly and sufficiently similar to past cases in which someone's steadfastness helped their epistemic community achieve its relevant epistemic goals*.

The advantage of (v) over (iv) and (iii), of course, is that S no longer requires a crystal ball. Now she needs her history books.³⁹

But it seems to me that Elgin could still reply that (v) is excessively demanding. Condition (v) requires S to make reasonable comparative judgments about her own case and past cases, and to reasonably conclude that hers is relevantly and sufficiently similar, and this might still be asking too much. For one thing, community inquiry might slow down too if S were required to reasonably establish what amounts to a legalistic precedent for her steadfastness prior to being entitled to it. For another, S's judgments about the relevance and sufficiency of the similarities between her case and past cases might themselves be controversial to the point where a higher-order problem looms: in order for S to reasonably remain steadfast about her first-order judgment that P, she would have to have second-order judgments about precedents concerning which she could reasonably remain steadfast – and doing so would require a third-order assessment of her epistemic situation with respect to these precedents. The mind begins to boggle. So, in short, if all this is plausible, then Elgin can respond that the latest proffered condition, (v), is, like its ancestors (iv) and (iii), sufficient but not necessary for licensing S's steadfastness.⁴⁰

6. DESIDERATA FOR FURTHER WORK

In the previous section, I set out and responded to seven objections to Elgin's community-oriented steadfastness. I take myself to have shown that none of these is decisive. Still, it must be said that Elgin has presented only a fairly brief sketch of her view to date. So, in this final section, I identify some fruitful ways in which her view could be developed further.

The final objection I considered was that nobody can be sure enough about whether their community will achieve its relevant epistemic goals because of their steadfastness about P for interval t . I fleshed this out in three ways. The first proposal for a necessary condition held that S must have reasonable acceptances about the *all-things-considered* consequences of her steadfastness. The second proposal held that S must have reasonable acceptances about the consequences of her steadfastness, *given the evidence she can reasonably be expected to acquire*. The third proposal held that S must be reasonable in thinking that her steadfastness is *relevantly and sufficiently similar to past cases* in which someone's steadfastness helped their epistemic community achieve its relevant epistemic goals.

Notice that the first two of these proposals concentrated on consequences of a certain *individual action*: S's act of remaining steadfast for interval t . Act-consequentialism, of course, is the ethical view that privileges the evaluation of the consequences of an *individual action* in assessing that action's moral status. Likewise, *epistemic* act-consequentialism privileges the evaluation of the consequences of an individual action in assessing that action's epistemic status. Act-consequentialism, however, is widely thought to be deeply implausible.⁴¹ So perhaps it's simply uncharitable to attempt to flesh out Elgin's view in simplistic act-consequentialist terms.

And yet, as I noted earlier, Elgin's view does seem thoroughly consequentialist: her defence of the propriety of steadfastness, after all, focusses precisely on the epistemically salutary consequences of steadfastness for epistemic communities. Rule-consequentialism, of course, has exactly this community-oriented focus. As Tim Mulgan explains, rule-consequentialism "... reflects a picture of morality as a task given, not to isolated individual rational agents ... but to a particular *community of human beings*" (2007, 124). Given this, I suggest that it would be better to understand Elgin's defence of steadfastness in epistemic *rule-consequentialist* terms.⁴²

For starters, here is an important statement of rule-consequentialism by one of its best-known contemporary proponents, Brad Hooker:

An act is wrong if and only if it is forbidden by the code of rules whose internalization by the overwhelming majority of everyone everywhere in each new generation has maximum expected value in terms of well-being ... (2000, 32).

This definition is easily translated into the language of permissions:

An act is permissible if and only if it is permitted by the code of rules whose internalization by the overwhelming majority of everyone everywhere in each new generation has maximum expected value in terms of well-being
...

And it's easy to cook up an epistemic analogue:

An epistemic act is rationally permissible if and only if it is permitted by the code of rules whose internalization by the overwhelming majority of everyone everywhere in each new generation has maximum expected epistemic value.

Now, as Dale Miller notes, rule-consequentialists needn't formulate their view with respect to a maximally large community, like Hooker does: "[a]nother possibility is to narrow the group in question, so that it includes only the members of a given society or, to narrow it even further, only those individuals who are members of a given society during a particular period of time" (2014, 152). This suggests the following modification:

An epistemic act is rationally permissible if and only if it is permitted by the code of rules whose internalization by the overwhelming majority of a community of inquiry has maximum expected epistemic value.

I do not intend to suggest that Elgin is, or must be, wedded to either of these last two forms of epistemic rule-consequentialism. Instead, my point is just that Elgin's view seems rather amendable to a rule-consequentialist framework, the details of which could be developed in various ways.

If this much is correct, then there is much work to be done. There are, of course, many forms of rule-consequentialism, and many arguments for and

against these views. A thorough assessment would involve examining and assessing these – and, crucially, their epistemic analogues.⁴³ While space does not permit such an undertaking, it will be useful to consider, in closing, what a steadfastness-licensing rule, embedded within an epistemic rule-consequentialist code, might look like.⁴⁴ Here are eight issues well worth considering in any such attempt:

- In Section 5, I noted that Elgin focusses on cases in which “... reasons favoring each side of a dispute are sparse or exceedingly delicate, or the evidence is equivocal, or each side can solve important common problems that the other cannot” (2010, 67). We can think of these as characteristics of the contexts of disagreements in which an epistemic rule could license steadfastness on consequentialist grounds. Further work should attempt to identify other such characteristics, and to develop all of them, in an effort to clarify the contexts in which the relevant steadfastness-permitting rules apply.
- In Section 5, I also noted that Elgin should restrict her focus to disagreements in which each side’s position surpasses some minimal epistemic threshold, in order to prevent her view from licensing steadfastness for absurd positions. Well-developed epistemic rules should ideally encode an account, at least in general terms, of what this threshold looks like. Nuance will be needed here. The threshold must be high enough to preclude entirely spurious views, while being low enough to permit, for example, highly counterintuitive scientific hypotheses that may well bear fruit.⁴⁵ It may also be that the threshold is context-dependent: perhaps some areas of inquiry can better tolerate, or are better served by, a lower threshold, while others require a higher one. Accordingly, it will be no easy task to develop such an account.
- More work is also needed to determine what counts as an ‘epistemic benefit’ for a community of inquiry. There may well be many different types of epistemic benefits, and which ones are relevant may well also be context-dependent. An Elgin-inspired rule that license steadfastness should flesh this out appropriately.
- ‘Epistemic benefits’ evidently come in degrees: some are large; others are small. Clearly, a rule that license steadfastness in the relevant contexts will not have significant expected epistemic utility if its internalization by the relevant group tends to reap only very minor benefits. Accordingly, it would be immensely helpful if the rule could encode an account of what degree of epistemic benefit, for a community, steadfastness ought to promote, at least in general terms.
- I noted earlier that one way to formulate epistemic rule-consequentialism is to relativize its rules to a community of inquiry. Perhaps this is more plausible than (the epistemic analogue) of Hooker’s rule-consequentialism, which stipulates that the ideal code of rules is internalized by “everyone everywhere”. However, doing so requires greater clarity about what is meant, for these purposes, by a ‘community of inquiry’. Presumably one marker of such a

community is that its members share certain epistemic goals, but questions immediately arise: What exactly is an epistemic community? How many goals must they share? Must they explicitly share them, or could they share them implicitly? Attention to questions such as these will help to clarify the appropriate contexts of application of the relevant rules.⁴⁶

- Relatedly, it seems obvious that many individuals belong to multiple epistemic communities. Suppose that S has multiple memberships, and suppose, further, that she reasonably thinks that it's best for one of her communities if S accepts P, and also best for another one of her communities if S accepts $\sim P$. What should she do? A more developed account of a steadfastness-licensing rule should address what rationality permits and what it forbids in such a situation.
- This paper has largely concentrated on Elgin's effort to establish the rational permissibility of steadfastness in certain types of disagreements. But, of course, to say that remaining steadfast is rationally permissible is not to say that conciliating is rationally impermissible. A full account of the relevant epistemic rules should indicate, at least in general terms, what other attitudes, besides accepting P (for the sake of one's community) are rationally permissible, and in what cases.⁴⁷
- A final desideratum is that the development of an epistemic rule-consequentialist framework be appropriately informed by empirical findings. As we saw in Section 4, empirical work does indeed suggest that in controlled settings, the viewpoint diversity licensed by Elgin's defence of steadfastness can indeed yield positive epistemic outcomes. That said, we should not be unduly restrictive here: the lab does not always resemble the real world, and I take it that we ultimately want to devise a code of epistemic rules for the latter.

CONCLUSION

After setting out the relevant dialectical context in Sections 1 and 2, and explaining Catherine Elgin's *community-oriented steadfastness* in Section 3, I argued in Section 4 that her view has considerable initial plausibility: it is supported by significant empirical evidence, it comports well with accounts of the division of cognitive labour suggested by philosophers of science, and it fits nicely with the evolutionary-adaptive account of confirmation bias suggested by the argumentative theory of reason. In Section 5, I defended Elgin's view against seven objections, arranged in increasing order of strength. Finally, in Section 6 I suggested that her view can be understood in terms of epistemic rule-consequentialism, and I outlined eight ways in which it could fruitfully be developed further. Catherine Elgin's *community-oriented steadfastness* is a creative and important response to the problem of disagreement, and I hope that philosophers will give it the attention it deserves, perhaps by developing it in the ways outlined here.

ACKNOWLEDGEMENTS

I presented ancestors of this paper at the *Cave Hill Philosophy Symposium*, University of the West Indies at Cave Hill (November 12th, 2014); at Ryerson University's 8th Annual Philosophy Symposium (May 15th, 2015); at the *Canadian Philosophical Association Conference*, University of Ottawa (June 1st, 2015); at the *Western Canadian Philosophical Association Conference*, Saskatoon (October 3rd, 2015); at a symposium on faith and doubt held at Tyndale University College, (March 10th, 2016); and at a symposium on disagreement, higher-order evidence, and new arguments for skepticism, held at the *Canadian Philosophical Association Conference*, Ryerson University (May 29th, 2017). I am grateful to all my interlocutors on those occasions. I am particularly grateful to Nathan Ballantyne, David Hunter, Kirk Lougheed, and two anonymous referees for very probing and helpful comments on earlier drafts.

NOTES

¹ Good surveys of this discussion can be found in Christensen (2009) and Lackey (2010).

² This view is defended in various ways by, for example, Feldman (2006), Christensen (2007), Elga (2007, 2010), and Ballantyne (2014).

³ This view is defended in various ways by, for example, Kelly (2005), Moffett (2007), Bergmann (2009), and Sosa (2010).

⁴ King's discussion avoids talk of degrees of belief, strictly for simplicity, but I have added it parenthetically here, in order to not foreclose the possibility that conciliationism can lead to diminished confidence that is nevertheless compatible with belief. I have also added the modifier 'epistemically', for reasons that will be discussed in Section 5.7.

⁵ Elgin adds: "The reason for the restriction to assertoric inference is to screen off premises used in *reductios*. The reason for the restriction to cognitive interests is that a premise accepted because it fosters noncognitive ends – because it is consoling or amusing, for example – is epistemologically irrelevant" (2010, 64). In a different article, she offers a similar definition of belief (2012, 81), but it becomes clear that what she really means is acceptance. And in yet another article, she refers to this stance as being *committed* (2018, 20). In what follows, I stick with the more familiar term 'acceptance'.

⁶ For other presentations of this argument, see Elgin (2012, 5-7; 2017, 111-117; and 2018, 18-19). For a case study from anthropology that is strikingly similar to Elgin's hypothetical example about paleontology, see De Cruz and De Smedt (2013).

⁷ Some of these studies are also discussed in Dunn (*ms.*); Matheson (2015a and 2015b); Lougheed (2018a and 2018b); and Hallson and Kappel (forthcoming). For some important caveats, see Hallson and Kappel (forthcoming, section 2.1).

⁸ Given all this, it's no wonder that a recent spate of middlebrow trade publications have made much of this research. See Surowiecki (2004); Page (2007); Howe (2008); and Weinberger (2011).

⁹ A good entry point into the vast philosophy of science literature on the division of cognitive labour is Muldoon (2013). In her 2010 and 2018 papers, Elgin cites a seminal paper in this area: Kitcher (1990).

¹⁰ I believe that Dunn (*ms.*) was the first philosopher in the disagreement debate to draw this connection to the argumentative theory of reason.

¹¹ Elgin, in fact, comes close to making the same point when she says that on her model, those who remain steadfast “have a strong incentive to strengthen and stabilize their position and to protect themselves should it turn out that they are wrong” (2018, 20).

¹² As I explain in Section 6, it will be extremely difficult to flesh out this minimal epistemic threshold precisely. But I don’t think that this presents a special challenge for Elgin’s view: everyone in the disagreement debate ultimately owes an account that specifies which sorts of views have enough going for them, epistemically speaking, to motivate conciliationism or steadfastness.

¹³ I should note, however, that this move relies on a different, enriched, account of acceptance. Elgin does not simply appeal here to her official definition of acceptance *simpliciter*, which I set out in Section 3. Rather, she has in mind the more specific idea of accepting a proposition *for the sake of one’s community* – and she emphasizes that *this*, properly construed, does not license stubbornness. Elgin is, of course, entitled to stipulate such an enriched account of what’s involved in acceptance in this context – but then the defender of steadfastness about belief can do likewise, which would undermine Elgin’s original complaint, set out in Section 2, that steadfastness leads to objectionable stubbornness. Be that as it may, I think that it is dialectically legitimate for Elgin to say that, properly understood, her view of acceptance *for the sake of one’s community* does not license stubbornness – and this is all she needs to defeat the objection. I thank Nathan Ballantyne and David Hunter for careful discussion of this issue, and an anonymous referee for pressing me to clarify it further.

¹⁴ Apparently this is a (very!) free-form paraphrase of *Nicomachean Ethics* I.1094b24 which actually states: “Each of our claims, then, ought to be accepted in the same way [as claiming to hold good, usually], since the educated person seeks exactness in each area to the extent that the nature of the subject allows” (*transl.* Irwin 1985).

¹⁵ A related proposal is due to Goldberg, who recommends adopting an attitude of *attitudinal speculation*: “[O]ne who attitudinally speculates that p regards p as more likely than not-p, though also regards the total evidence as stopping short of warranting belief in p” (2013, 283). See Barnett (2019) for discussion.

¹⁶ And see also Greitemeyer *et al.* (2006). These studies are also discussed in Dunn (*ms.*), Matheson (2015, 144–146), Lougheed (2018b), and Hallsson and Kappel (forthcoming).

¹⁷ I thank an anonymous referee for suggesting these two points.

¹⁸ Feldman (2006) makes this point *en passant*; Elgin (2010, 56) mentions it, and Ballantyne (2014) argues for it at length.

¹⁹ As Thomas Kelly puts the point: “... whether there is any actual disagreement might very well depend on factors that everyone will immediately recognize as irrelevant to the truth of the question at issue. (Suppose, for example, that there would be considerable

disagreement with respect to some issue, but that all of the would-be dissenters have been put to death by an evil and intolerant tyrant.)” (2005, 181).

²⁰ I am persuaded that the problem of possible peer disagreement is every bit as serious as the problem of actual peer disagreement, but I will not argue for this here.

²¹ I am grateful to Kirk Lougheed for pushing me to think this through more carefully.

²² That’s because it introduces nothing that Elgin has not already conceded.

²³ I have replaced ‘belief’ with ‘acceptance’ in this quotation, in order to remain consistent with Elgin’s usage. For similar charges, see Feldman (2006, 221 and 2011, 157).

²⁴ I thank an anonymous referee for pressing me to develop responses to the second and third of these.

²⁵ For a survey of various presentations of objections in this vein, see Lougheed and Simpson (2017). See also Dunn (*ms.*) for discussion.

²⁶ Dunn (*ms.*), Talbot (2014), Everett (2015), Matheson (2015a and 2015b), Lougheed and Simpson (2017), Fleischer (2018), Barnett (2019), and Hallsson and Kappel (forthcoming).

²⁷ They write: “... when an agent identifies some spectrum of rationally permissible doxastic attitudes in response to some body of evidence, and where she has some ability to influence what doxastic attitude she holds among that range of options, then, to the extent that she has reason to believe that her holding doxastic attitude X will have better global epistemic consequences than her holding doxastic attitude Y, she has a reason – one that may be aptly characterized as an epistemic reason – to favour X over Y (2017, 157).

²⁸ Leite is not himself a committed instrumentalist (2007, 554).

²⁹ I should note that Kelly (2003) has other objections, and also that Kelly (2007) replies to Leite (2007). Space does not permit further appraisal of this exchange.

³⁰ Space does not permit an extended discussion of (a) and (b). I should note that insofar as (b) is a version of *evidentialism*, the most comprehensive defence of this view is Conee and Feldmann (2004). The most prominent rival to this view is typically thought to be *reliabilism*. See, for example, Goldman (1986).

³¹ Here is why I hedge by saying ‘virtually’. A well-known objection to conciliationism holds that this view is self-defeating. (For discussion of this objection, see, for example, Elga 2010). Given that this is an objection to conciliationism, it might be thought to be an indirect defence of steadfastness. But clearly it carries no commitment to the denial of (a) and (b).

³² Moreover, the steadfast view is one way to defend a broader thesis in epistemology known as *permissivism* against its rival, *uniqueness*. For more on this debate, see Kopec, M. and Titelbaum, M. (2016) and Rosa (2018).

³³ Interestingly, Dunn and Ahlstrom-Vij (2018) urge that even *evidentialism* ought to be deemed a species of epistemic consequentialism.

³⁴ Although (ii) might appear to add nothing important to (i), I think it does: after all, it's easy to imagine that a view could satisfy all of the disjuncts listed in (i), while nevertheless failing to meet a plausible overall epistemic threshold.

³⁵ Lougheed (2018a) defends what he calls the "Benefits to Inquiry Argument", which is similar in spirit to Elgin's. He also considers an objection in this vein. But he dismisses it too hastily when he says: "This objection is about the *practical application* of the argument not the *truth* of any of the premises" (2018, 275). After all, the third premise in his own argument is this:

S reasonably believes that there are future epistemic benefits [to be gained] from continuing to believe *P* within the context of *R* in the face of peer disagreement about *P* (2018, 266).

If, as the objection posits, nobody can reasonably believe such a thing, then this does seem to count against the truth of this premise: it is true of no subject *S*.

³⁶ More precisely: sufficient to license *S*'s steadfastness *provided that* (i) and (ii) are also satisfied. I make the same presumption with respect to (iv) and (v) below.

³⁷ There may, of course, be some context-sensitivity here: perhaps it is reasonable to expect more evidence-gathering from *S* in some situations than in others. (Thanks to Nathan Ballantyne for pointing this out.)

³⁸ In fairness, one might wonder whether Semmelweis really satisfies the entire antecedent of (SK1*). But let's stipulate that, given his knowledge and environment, and the knowledge and environment of those who disagreed with him, there were at least some temporal slices of Semmelweis that do satisfy this antecedent.

³⁹ Of course not everyone is a Semmelweis, or, for that matter, a Copernicus or a Darwin. Nor, indeed, is every context in which Elgin wants to license steadfastness is relevantly and sufficiently similar to these cases. That, however, is not a problem, for (v) does not demand this. There might well be *other* cases, to which *S*'s situation is relevantly and sufficiently similar, that do not involve such intellectual giants and heroes.

⁴⁰ At this point in the dialectic, the burden of proof remains on the objector to specify a plausible necessary condition for *S* to be sufficiently confident that her remaining steadfast concerning *P* for time interval *t* will help her community to achieve its epistemic goals. There is a more ambitious way to respond to this objection. Instead of conceding that a necessary condition along these lines needs to be satisfied, and waiting for the objector to provide it, Elgin could say that, given the epistemic benefits of viewpoint diversity, there is no such condition. Instead, she could say, *S* should simply be *presumed* to be within her rights to remain steadfast about *P*, unless this presumption is defeated by, for example, concrete reasons to suppose that her steadfastness will not help the community in this particular case. I thank an anonymous referee for this suggestion.

⁴¹ And, of course, one well-worn reason for this is just what we have seen: it appears to demand that agents know far more about the consequences of their actions than is reasonable to expect.

⁴² I thank Gary Toop for originally pushing me to think this way. Lougheed and Simpson (2017, 159-160) also appear sympathetic to a rule-consequentialist defence of steadfastness

on community-oriented grounds. Relatedly, Leite also suggests that *instrumentalism* can be developed in *rule-utilitarian* terms (2007, 461).

⁴³ An excellent entry point into the literature on epistemic consequentialism is Ahlstrom-Vij and Dunn (2018).

⁴⁴ Moreover, I submit that these are useful directions for further inquiry even if Elgin's community-oriented defence of steadfastness is not best construed in rule-consequentialist terms.

⁴⁵ I thank an anonymous referee for suggesting the latter point.

⁴⁶ See Elgin (2017, 112-113) for some related discussion of what separates epistemic communities from other communities.

⁴⁷ In particular, conciliating for *community-oriented* reasons might be rationally permissible. Indeed, in some cases it might be rationally required.

REFERENCES

- Ahlstrom-Vij, H. and Dunn, J. (2018) *Epistemic Consequentialism*. Oxford. Oxford University Press.
- Ballantyne, N. (2014) “Counterfactual Philosophers.” *Philosophy and Phenomenological Research* 88: 368-387.
- Barnett, Z. (2019) “Philosophy Without Belief.” *Mind* 128: 109-138.
- Bergmann, M. (2009) “Rational Disagreement after Full Disclosure.” *Episteme: A Journal of Social Epistemology* 6: 336-353.
- Brodbeck, F. C., Kerschreiter, R., Mojzisch, A., Frey, D., and Schulz-Hardt, S. (2002) “The Dissemination of Critical, Unshared Information in Decision Making Groups: The Effects of Pre-Discussion Dissent.” *European Journal of Social Psychology* 32: 35-56.
- Christensen, D. (2007) “Epistemology of Disagreement: The Good News.” *Philosophical Review* 116: 187-217.
- (2009) “Disagreement as Evidence: The Epistemology of Controversy.” *Philosophy Compass* 4: 756-767.
- Cohen, L.J. (1992) *An Essay on Belief and Acceptance*. Oxford: Oxford University Press.
- Conee, E. and Feldman, R. (2004) *Evidentialism: Essays in Epistemology*. Oxford: Oxford University Press.
- Crisp, R.J. and Turner, R.N. (2011) “Cognitive Adaptation to the Experience of Social and Cultural Diversity.” *Psychological Bulletin* 137: 242-266.
- De Cruz, H. and De Smedt, J. (2013) “The Value of Epistemic Disagreement in Scientific Practice. The Case of *Homo Floresiensis*.” *Studies in History and Philosophy of Science* 44: 169-177.
- Duarte, J. L., Crawford, J. T., Stern, C., Haidt, J., Jussim, L., and Tetlock, P. E. (2015) “Political Diversity will Improve Social Psychological Science.” *Behavioral and Brain Sciences* 38: e130.
- Dunn, J. (ms.) “Peer Disagreement and Group Inquiry.”
- Dunn, J. and Ahlstrom-Vij, H. (2018) “Introduction.” In *Epistemic Consequentialism*. Edited by H. Ahlstrom-Vij and J. Dunn, 2-20. Oxford. Oxford University Press.
- Elga, A. (2007) “Reflection and Disagreement.” *Noûs* 41: 478-502.

- _____. (2010) "How to Disagree about How to Disagree." In *Disagreement*. Edited by R. Feldman and T. Warfield, 175-186. Oxford: Oxford University Press.
- Elgin, C. (2010) "Persistent Disagreement." In *Disagreement*. Edited by R. Feldman and T. Warfield, 53–68. Oxford: Oxford University Press.
- _____. (2012) "Begging to Differ." *The Philosopher's Magazine*. Fourth Quarter: 77-82.
- _____. (2017) *True Enough*. Cambridge: MIT Press.
- _____. (2018) "Reasonable Disagreement." In *Voicing Dissent: The Ethics and Epistemology of Making Disagreement Public*. Edited by C. R. Johnson, 10-21. New York: Routledge.
- Everett, T. J. (2015) "Peer Disagreement and Two Principles of Rational Belief." *Australasian Journal of Philosophy* 93: 273-286.
- Feldman, R. (2006) "Epistemological Puzzles about Disagreement." In *Epistemology Futures*. Edited by S. Hetherington, 216-236. Oxford: Oxford University Press.
- _____. (2011) "Reasonable Religious Disagreements." In *Social Epistemology: Essential Readings*. Edited by A. I. Goldman and D. Whitcomb, 137-157. Oxford: Oxford University Press.
- Fleischer, W. (2018) "Rational Endorsement." *Philosophical Studies* 175: 2649-2675.
- Gigone, D. and Hastie, R. (1997) "Proper Analysis of the Accuracy of Group Judgments." *Psychological Bulletin* 121: 149-167.
- Goldberg, S. (2013) "Defending Philosophy in the Face of Systematic Disagreement." In *Disagreement and Skepticism*. Edited by D.E. Machuca, 277-294. New York: Routledge.
- Goldman, A. (1986) *Epistemology and Cognition*. Cambridge: Harvard University Press.
- Greitemeyer, T., Schulz-Hardt, S., Brodbeck, F. C., and Frey, D. (2006) "Information Sampling and Group Decision Making: The Effects of an Advocacy Decision Procedure and Task Experience." *Journal of Experimental Psychology: Applied* 12: 31-42.
- Hallsson, B. and Kappel, K. (forthcoming) "Disagreement and the Division of Epistemic Labor." *Synthese*.

- Hastie, R. (1986) "Experimental Evidence on Group Accuracy." In *Information Pooling and Group Decision Making: Proceedings of the Second University of California, Irvine Conference on Political Economy*. Edited by B. Grofman and G. Owen, 129-157. JAI Press.
- Hooker, B. (2000) *Ideal Code, Real World: A Rule-Consequentialist Theory of Morality*. Oxford: Oxford University Press.
- Howe, J. (2008) *Crowdsourcing: Why the Power of the Crowd Is Driving the Future of Business*. New York: Three Rivers Press.
- Irwin, T., Transl. (1985) *Aristotle: Nicomachean Ethics*. Indianapolis: Hackett Publishing Company.
- Jehn, K. A., Northcraft, G. B., and Neale, M. A. (1999) "Why Differences Make a Difference: A Field Study of Diversity, Conflict, and Performance in Workgroups." *Administrative Science Quarterly* 44: 741-763.
- Kelly, T. (2003) "Epistemic Rationality as Instrumental Rationality: A Critique." *Philosophy and Phenomenological Research* 66: 612-640.
- _____. (2005) "The Epistemic Significance of Disagreement." In *Oxford Studies in Epistemology, Volume 1*. Edited by J. Hawthorne and T. Szabó Gendler, 167-196. Oxford: Oxford University Press.
- _____. (2007) "Evidence and Normativity: Reply to Leite." *Philosophy and Phenomenological Research* 75: 465-474.
- King, N. (2012) "Disagreement: What's the Problem? Or A Good Peer is Hard to Find." *Philosophy and Phenomenological Research* 85: 249-272.
- Kitcher, P. (1990) "The Division of Cognitive Labour." *Journal of Philosophy* 87: 5-22.
- Kopiec, M. and Titelbaum, M. (2016) "The Uniqueness Thesis." *Philosophy Compass* 11: 189-2000.
- Krause, S., James, R., Faria, J. J., Ruxton, G. D., and Krause, J. (2011) "Swarm Intelligence in Humans: Diversity can Trump Ability." *Animal Behaviour* 81: 941-948.
- Kuhn, D., Shaw, V. and Felton, M. (1997) "Effects of Dyadic Interaction on Argumentive Reasoning." *Cognition and Instruction* 15: 287-315.
- Lackey, J. (2010) "The Epistemology of Disagreement." *Oxford Bibliographies Online*.

- Lamm, H. and Myers, D. G. (1978) "Group-Induced Polarization of Attitudes and Behavior." *Advances in Experimental Social Psychology* 11: 145-195.
- Laughlin, P. R. and Ellis, A. L. (1986) "Demonstrability and Social Combination Processes on Mathematical Intellective Tasks." *Journal of Experimental Social Psychology* 22: 177-189.
- Leite, A. (2007) "Epistemic Instrumentalism and Reasons for Belief: A Reply to Tom Kelly's 'Epistemic Rationality as Instrumental Rationality: A Critique.'" *Philosophy and Phenomenological Research* 75: 456-464.
- Lougheed, K. (2018a) "The Epistemic Value of Deep Disagreements." *Informal Logic* 38: 263-292.
- _____. (2018b) "Recognition and Epistemic Injustice in the Epistemology of Disagreement." *Philosophical Forum* 49: 363-377.
- Lougheed, K. and Simpson, R. (2017) "Indirect Epistemic Reasons and Religious Belief." *Religious Studies* 53: 151-159.
- Matheson J. (2015a) *The Epistemic Significance of Disagreement*. New York: Palgrave Macmillan.
- _____. (2015b) "Disagreement and the Ethics of Belief." In *The Future of Social Epistemology: a Collective Vision*. Edited by J. H Collier, 139-148. USA: Rowman and Littlefield.
- Mercier, H. and Sperber, D. (2011) "Why do Humans Reason? Arguments for an Argumentative Theory." *Behavioural and Brain Sciences* 34: 57-111.
- _____. (2017) *The Enigma of Reason*. Cambridge: Harvard University Press.
- Miller, D. (2014) "Rule Utilitarianism." In *The Cambridge Companion to Utilitarianism*. Edited by B. Eggleston and D. Miller, 146-165. Cambridge: Cambridge University Press.
- Moffett, M. (2007) "Reasonable Disagreement and Rational Group Inquiry." *Episteme: A Journal of Social Epistemology* 4: 352-367.
- Moscovici, C. and Personnaz, B. (1980) "Minority Influence and Conversion Behavior in a Perceptual Task." *Journal of Experimental Social Psychology* 16: 270-282.
- Moscovici, S., and Zavalloni, M. (1969) "The Group as a Polarizer of Attitudes." *Journal of Personality and Social Psychology* 12: 125-135.
- Moshman, D. and Geil, M. (1998) "Collaborative Reasoning: Evidence for Collective Rationality." *Thinking and Reasoning* 4: 231-48.

- Muldoon, R. (2013) "Diversity and the Division of Cognitive Labour." *Philosophy Compass* 8: 117-125.
- Mulgan, T. (2007) *Understanding Utilitarianism*. Stocksfield: Acumen Publishing Limited.
- Nemeth, C. J. (1986) "Differential Contributions of Majority and Minority Influence." *Psychological Review* 93: 23-32.
- Nemeth, C. J. (1995) "Dissent as Driving Cognition, Attitudes, and Judgments." *Social Cognition* 13: 273-291.
- Nemeth, C., Brown, K., and Rogers, J. (2001) "Devil's Advocate Versus Authentic Dissent: Stimulating Quantity and Quality." *European Journal of Social Psychology* 31: 707-720.
- Page, S. (2007) *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*. Princeton: Princeton University Press.
- Rosa, L. (2018) "Uniqueness and Permissiveness in Epistemology." *Oxford Bibliographies Online*.
- Schulz-Hardt, S., Brodbeck, F. C., Mojzisch, A., Kerschreiter, R., and Frey, D. (2006) "Group Decision Making in Hidden Profile Situations: Dissent as a Facilitator for Decision Quality." *Journal of Personality and Social Psychology* 91: 1080-1093.
- Schulz-Hardt, S., Jochims, M. and Frey, D. (2002) "Productive Conflict in Group Decision Making: Genuine and Contrived Dissent as Strategies to Counteract Biased Information Seeking." *Organizational Behavior and Human Decision Processes* 88: 563-86.
- Sosa, E. (2010) "The Epistemology of Disagreement." In *Social Epistemology*. Edited by A. Haddock, A. Millar, and D. Pritchard, 278–297. Oxford: Oxford University Press. [First published in *Episteme: A Journal of Social Epistemology* 6.3 (2009): 269–279.]
- Sperber, D. (2001) "An Evolutionary Perspective on Testimony and Argumentation." *Philosophical Topics* 29: 401-413.
- Sunstein, C. (2002) "The Law of Group Polarization." *Journal of Political Philosophy* 10: 175-95.
- Surowiecki, J. (2004) *The Wisdom of Crowds: How the Many Are Smarter than the Few and How Collective Wisdom Shapes Business, Economies, Societies and Nations*. New York: Doubleday.

- Talbot, B. (2014) "Truth-Promoting Non-Evidential Reasons for Belief." *Philosophical Studies* 168: 599-618.
- Trouche, E., Sander, E., and Mercier, H. (2014) "Arguments, More than Confidence, Explain the Good Performance of Reasoning Groups." *Journal of Experimental Psychology: General* 143: 1958-1971.
- van Knippenberg, D. and Schippers, M. C. (2007) "Work Group Diversity." *Annual Review of Psychology* 58: 515-541.
- Vinokur, A., and Burnstein, E. (1978) "Depolarization of Attitudes in Groups." *Journal of Personality and Social Psychology* 36: 872-885.
- Wienberger, D. (2011) *Too Big to Know: Rethinking Knowledge Now That the Facts Aren't the Facts, Experts Are Everywhere, and the Smartest Person in the Room Is the Room*. New York: Basic Books.
- Woolley, A. W., Aggarwal, I., and Malone, T.W. (2015) "Collective Intelligence and Group Performance." *Current Directions in Psychological Science* 24: 420-424.