

Forthcoming in *Voicing Dissent*, ed. Casey Johnson. Routledge.

Reasonable Disagreement¹

Catherine Z. Elgin

Normally, we take it that we know rather a lot. We take ourselves to know the date of the Battle of Trafalgar (1805), the atomic number of gold (79), the score of last week's football game (Arsenal beat Manchester United, 2-0), the name of the Queen of Spain (Sofia). We have reasons -- often quite good reasons -- for believing the things we take ourselves to know. Evidence mounts up, eventually reaching or surpassing the threshold required for knowledge. But typically our reasons are less than conclusive. They just render our conclusions highly probable. Moreover, some of the reasons may themselves be epistemically precarious. Perhaps some of the evidence is misleading or unreliable. Perhaps a source we trusted was misinformed. Perhaps an inference was unsound. Given our epistemic vulnerability, it is no surprise that others disagree with us about many of the things we take ourselves to know. What should we make of that? Typically philosophers assume that where parties disagree about a matter of fact at least one of them has made a mistake. I will argue that that is false. Because judgment involves choices from within ranges of epistemically acceptable alternatives, not all disagreements, and not all peer disagreements, about matters of fact rest on mistakes. This is not to say that the conclusions of both parties are true. It is to say neither party made an error in coming to her conclusion. If this is right, it reveals something important about the precariousness of our epistemic condition.

The standard question with respect to disagreement is how epistemic agents should respond. Should we hold fast to our opinions, revise or moderate them in the face of disagreement, jump to the other side? When the question is framed so schematically, the alternatives seem stark:

- A. Be steadfast: Hold on to your convictions, disregarding disagreement.
- B. Conciliate: Moderate your conviction in the face of disagreement. Either suspend judgment about the issue under dispute or lower your credence in your conclusion.
- C. Concede: Endorse your opponent's position, abandoning your own.

We need not think that the same stance is appropriate in every case. Some disagreements may call for steadfastness, others for conciliation. In yet others, the best thing to do is concede that your opponent is right. Let's look at some cases.

Hal is and knows himself to be pretty clueless about ornithology and rightly considers Val quite knowledgeable. He thinks that all goldfinches look alike. Val disagrees. She thinks goldfinches display sexual dimorphism. In the face of their disagreement, Hal probably ought to concede that she is right, or at least suspend judgment. He recognizes that she has expertise that he lacks. Their situations are not symmetrical. The disagreement gives Val no grounds for revising her opinion. Since Hal can hardly tell a sparrow from a hawk, his opinion need not give her pause. It is reasonable for her to remain steadfast. *Ceteris paribus*, when an expert and a novice disagree, the novice should back off. He should then either conciliate or concede. The expert, however, can reasonably hold fast.

What if parties to a disagreement are epistemic peers? Neither is more knowledgeable than the other. Here their level of competence matters. Bill and Jill are on the verge of failing introductory chemistry. Jill thinks that methane is an organic chemical; Bill thinks that it is

inorganic. Given the precariousness of their grasp of the topic, they should probably suspend judgment and look it up. If an agent's grasp of an issue is sufficiently weak, it should be unsettled by the news that anyone, even an equally ignorant peer, disagrees.

Cases like these are easily handled. Disagreements among competent epistemic peers are more troublesome. What should a responsible, competent epistemic agent do on learning that a peer with the same evidence, background assumptions, reasoning abilities, and epistemic incentives disagrees?² In what follows I will assume that the peers under discussion are competent with regard to the issue in dispute. Concessiveness is reasonable only when an epistemic agent believes that her interlocutor has considerably greater relevant expertise than she does. Where peers disagree, there is no reason to be concessive. In peer disagreements, an epistemic agent should either stick to her guns, or she should conciliate. Kelly (2005) advocates steadfastness. If Nell considers herself to have been responsible in forming her opinion, she should dismiss her peer Mel's opinion, concluding that he must be in error. Christensen (2007) advocates adopting a conciliatory stance. Nell should lower her credence or suspend judgment, concluding that one of the two is in error. Either way, it is widely agreed that at least one has made a mistake -- miscalculated, overlooked evidence, ignored base rates, reasoned fallaciously, or committed some other cognitive gaffe. Whatever the precise failing, where peers disagree, at least one of them did something cognitively culpable. As Sidgwick says, 'If I find any of my judgments, intuitive or inferential, in direct conflict with some other mind, there must be error somewhere' (1981 p. 342).

Some disagreements, such as Christensen's restaurant case, exhibit this profile. When competent agents who are equally good at mental math disagree about how to split the bill, at least one of them miscalculated. The obvious remedy is to suspend judgment temporarily, do the

calculation publicly, and figure out the right answer (2007). Similarly if they disagree about the date of the Battle of Trafalgar. That too can be easily checked. Pretty plainly, parties to the dispute should suspend judgment, pull out their smart phones and google the answer. Both questions are easily resolved. To suspend judgment and check the answer takes little time and effort. Conciliation in such cases is relatively cost free.

But many disagreements lack an easy resolution. There is no feasible way to backtrack and find an error. Sometimes the information that would settle the dispute is simply not to be had. Ann and Dan, forty-year-olds with good memories, disagree about which team won the third grade sack race when they were in school. No records were kept about this momentous event. Nor are they in touch with any of their former classmates. That being the case, their disagreement seems fated to remain unresolved. In cases like this, the protagonists agree about everything relevant except the fact at the heart of the dispute. Someone made an error, but there is no way to determine who.

In other cases, however, there is no error. That is not to say that both claims are true. It is to say that both opinions are products of epistemically acceptable, perhaps even epistemically impeccable, reasoning. Although one of the protagonists arrived at a false conclusion, neither made a mistake. I want to focus on such a case.

Suppose an animal bone with holes pierced in it is found among the artifacts in a Neanderthal settlement.³ It looks enough like a primitive flute that had it been found in a prehistoric Homo sapiens settlement, paleoanthropologists would have no reservations about calling it a flute. But according to current theory, Neanderthal brains were not complex enough to create music, nor were their hands dexterous enough to finger a flute. Jen and Ken, both eminent paleoanthropologists, disagree about whether the artifact is a Neanderthal flute. Jen

thinks it is a Neanderthal flute; Ken thinks it is not. Their disagreement may be irresolvable. Given the paucity of evidence about Neanderthals, there is no expectation that anyone will ever come up with sufficient evidence to settle the matter conclusively. Moreover, something significant is at stake here. The verdict as to whether the artifact is a Neanderthal flute will contribute to the ongoing debate about how and in what respects Neanderthals differed from early Homo sapiens. If Neanderthals were capable of making music, a host of related hypotheses are seriously off the mark.

Jen and Ken are equally knowledgeable and competent, and reach or surpass the threshold for being qualified to judge the matter under dispute. They have the same evidence and the same ability to assess the evidence. Both consider the issue significant. Neither is cavalier. What should they do? Conceding has no appeal. Neither party to this dispute has reason to think that the other is in a better position to assess the matter. But both conciliation and steadfastness look like live options. What is to be said for them?

The case for steadfastness: Jen knows that she has considered the matter carefully and judiciously and has taken account of all the available evidence. She knows that she would have no reservations about her conclusion if she relied only on the first-order evidence, her background assumptions, and the inferences she draws from them. She rightly considers Ken her intellectual peer in paleoanthropology. So she is surprised that Ken does not agree with her. If she is steadfast, she does not, however, consider his disagreement a reason to rethink her position or moderate her confidence level. She concludes that in this case Ken must have made a mistake. Maybe he was careless in his reasoning; maybe he overlooked some relevant factor. If he had thought the matter through as carefully as she did, she thinks, he would have come to the same conclusion. Ken's disagreement gives her no reason to change her mind (see Kelly 2005).

Ken's situation is symmetrical. He too recognizes that he thought the matter through carefully and judiciously. Being equally steadfast, he can only conclude that Jen, although normally quite solid on such issues, has made a regrettable mistake. Her disagreement gives him no reason to change his view either. If they remain steadfast, Jen and Ken may permanently be at loggerheads. Since each is inclined to dismiss the other's opinion, and to have no incentive to examine the other's reasons, neither is in a position to learn from the other. The disagreement gives them no reason to rethink the issue. The steadfast are *intellectually arrogant*. Faced with a disagreement, a steadfast epistemic agent immediately concludes that the other party must be wrong. They are *dogmatic*. They do not think that a peer's disagreeing gives them any reason to rethink their position or change their mind. They are, moreover, *disrespectful* of those they consider intellectual equals. Indeed, we might wonder whether they really *do* consider their opponents intellectual equals, if they are so quick to claim superiority the moment a disagreement arises.

The case for conciliation: We've seen that Jen has thought through the matter carefully and judiciously. She would have no reservations about her conclusion if she relied exclusively on the first-order evidence, her background assumptions and the inferences she draws from them. But Ken, whom she considers her epistemic peer, disagrees with her. She has no reason to think that Ken was any less careful and judicious in his approach to the question, or that he was cavalier in his reasoning or interpretation of the evidence. One of them must be wrong. But given that they are peers, there is currently no way to tell which one. She therefore suspends judgment, hoping that further evidence will be found that will decide the issue between them. Conciliators are *intellectually humble*. Their epistemic stance is *fallibilist*. Jen thinks she is right, but acknowledges that she could be wrong. Conciliators are *respectful* of their epistemic

peers. As far as Jen can tell, Ken's view has as much going for it as hers. There is no basis for choosing between them.

It might seem that this shows that when epistemic peers disagree, they should conciliate. Certainly conciliators seem more congenial than the steadfast. But conciliation can be costly. If Jen and Ken take their disagreement as a reason to suspend judgment about whether the artifact is a Neanderthal flute, neither of them is in a position to know. They have backed off from belief about the matter. If Jen's first-order evidence was actually sufficient for her to know, then by suspending judgment she sacrifices knowledge. For knowledge requires belief. One of them is surely right. The artifact either is a Neanderthal flute or it is not. If both hold fast to their opinions, then one of them has a true belief. And if the one with the true belief has sufficient evidence that bears non-accidentally on its truth, one of them knows. If Jen is right and sticks to her guns, she knows or at least justifiably believes that the object is a Neanderthal flute. So, it seems, the steadfast are capable of preserving knowledge while the conciliatory willingly give it up.

Steadfastness, however, instantiates Kripke's so-called dogmatism paradox (2011).

If S knows that p , then p is true.

If p is true, then any evidence against p is misleading.

If S wants to retain her knowledge, she should disregard misleading information.

So S should be dogmatic and disregard any evidence that seems to tell against p .

This argument is perfectly general. It applies to first-order evidence as well as evidence that S might glean from the discovery that T disagrees with her. It is not strictly a paradox, but it yields an epistemically disheartening recommendation. Rationality requires that epistemic agents be open-minded; they should not intentionally blind themselves to evidence, even if that evidence

might turn out to be misleading. So if Ken's disagreement is evidence that Jen is wrong, Jen should not disregard it. Although the steadfast can retain knowledge or highly justified belief, they do so by shirking their epistemic responsibility. It is just by luck that the evidence that might be gleaned from taking the disagreement seriously turned out to be misleading. Given that Jen and Ken are epistemic peers, it is just by luck that Jen (rather than Ken), by being steadfast, retains her knowledge.

Conciliators, however, might seem spineless. The moment they learn that a peer disagrees, they back off from full belief. It seems epistemically ill-advised, and perhaps a bit cowardly, to move too quickly from 'I might be wrong' to 'I should immediately abandon my position'. If this is what conciliation requires, it is not an attractive option. There is, however, some slack in the conciliationist position. Conciliators hold that on learning that a peer disagrees, an epistemic agent should either *suspend judgment* or *lower her credence*. Suspending judgment is a lot to ask, at least if that means having no opinion or thinking one is unjustified in one's opinion about a matter. Given the range of disagreement, most of us would have to suspend judgment about a lot. But it may be reasonable to say that we should lower our credence in our conclusion. Conciliators could still be fairly confident that their conclusion holds, but they would have to admit that they are not fully confident and are not entitled to be fully confident about it. This might result in a loss of knowledge. Whether it does depends on where the threshold for knowledge is, and how far above the threshold an agent's belief originally was. Arguably, it often *should* result in a loss of knowledge. The agent would not think that she is, or should consider herself to be, utterly clueless about the topic or spineless in backing off. But she would and should have her confidence in her conviction shaken. She might then still be pretty sure of the conclusion.

Conciliators seem to gain more than they lose. We saw that the steadfast are dogmatic. They dismiss their opponent's opinion as mistaken without further investigation. If steadfast Ken is wrong, he is likely to remain wrong, since he takes himself to have no reason to rethink the issue. The existence of the disagreement gives him no incentive to investigate the matter further. Conciliators have such an incentive. Jen and Ken have the same evidence and the same background information. On learning that Ken disagrees with her, Jen has reason to ask, 'What does he see that I do not?' She has an incentive to go back and examine the evidence, the background assumptions, her reasoning, and, if she has access to it, his. Has she overlooked something? Has she reasoned fallaciously? Is there an illuminating perspective on the issue that she missed? Conciliators are in a better position than the steadfast to find the mistake that led to the disagreement, assuming of course, that there was a mistake.

But the assumption that disagreement always rest on a mistake itself is a mistake. The issue is why the peers disagree. Even in the idealized scenario, where peers are identical in multiple relevant respects, there can be a variety of reasons. Although it is stipulated that the peers have the same evidence, background assumptions, reasoning abilities, and epistemic motivation, there is no reason to assume that they use them in the same way. If they do not, they may arrive at different verdicts. And there may be nothing clearly mistaken about the way either of them proceeds. In such cases, what is at issue is not just whether p , but how one ought to determine whether p . I will urge that our epistemic situation is more complex than is usually recognized. Peer disagreements bring this complexity to the fore.

Disagreements occur in a multi-dimensional space. Along a variety of axes, epistemic requirements fix ranges within which acceptable verdicts must lie. But the requirements are not sufficiently fine-grained to assure uniqueness or to provide a decision procedure for

differentiating among judgments that fall within the range of acceptability. If parties disagree because they make different choices within a range of epistemically acceptable options, neither has made a mistake. Where an agent comes down depends to a considerable extent on which alternatives she had chosen. Had she made other choices within the acceptable range, she would have reached a different verdict. Epistemic peers are not epistemic clones.

By stipulation, epistemic peers have the same evidence. But it does not follow that they assign the same weight to different bits of evidence. The disputed artifact was found in a cave in Slovenia, in an area known to have been populated by Neanderthals. Jen considers this highly significant, thinking it constitutes fairly strong evidence that the artifact was crafted by a Neanderthal. Ken attaches less weight to this datum, noting that the perforated bone could have been left in the cave by a wandering Homo sapiens, or been brought there by a Neanderthal, having been crafted elsewhere by a Homo sapiens. Ken does not deny that the location where the object was found is evidence, he just considers it relatively weak evidence.

Peers are said to have the same reasoning abilities. But they need not have the same reasoning styles. Ken tends to reason analogically and to credit analogical arguments. Jen, although equally adept at analogical reasoning, prefers inference to the best explanation, considering analogical arguments rather loose. Ken saw his dog pick up a ball in the yard this morning and carry it over to her bed. He buttresses his belief that the artifact might have been transported to the Neanderthal site by relying on the analogy with his dog's behavior. Jen thinks that the best -- simplest, most plausible -- explanation of an artifact's being found in a Neanderthal site is that it was crafted by a denizen of that site -- viz., a Neanderthal. Ken is skeptical of inference to the best explanation, thinking that such linear reasoning tends to blind one to available alternatives.

In characterizing a newly found object as a Neanderthal artifact, paleoanthropologists compare it with objects that are already characterized as such. If it is close enough to acknowledged Neanderthal artifacts in relevant respects, it is likely to be incorporated into the precedent class; otherwise it is apt to be rejected. But everyone recognizes that the current precedent class contains some items that were wrongly classified. That is, it contains misleaders. Ken and Jen differ over which items are misleaders. This leads them to different verdicts about this artifact. Jen credits scored bits of flint that are thought to have been used as tools. She thinks that if the Neanderthals had the cognitive and physical dexterity to carve such tools, they had the cognitive and physical dexterity to pierce holes in a bone to make a flute. Ken doubts that the gashes in the flint were intentionally produced. He suspects that what Jen considers intentional scores are in fact natural striations, amplified over eons by wind and water. So the similarity to the flint tools leads Jen to think the perforated bone is a Neanderthal artifact, while Ken harbors doubts. It may be obvious that an evidence class contains misleaders. Perhaps various elements are in tension with one another in the sense that if one is a Neanderthal artifact another probably is not. But it may not be obvious which bits of evidence are the misleaders. Peers then may disagree because they differ over which members of the (admittedly somewhat messy) evidence class they consider worthy of trust.

A subject's background information about a topic consists of all the information she has that directly or indirectly bears on that topic. Although epistemic peers by stipulation have the same background information, they need not draw on the same bits of information or assign the bits they draw on the same weight. If they do not, background information plays different roles in their reasoning. Contemporary anatomical theory suggests, albeit weakly, that organisms with thick phalanges are relatively lacking in fine motor control. Because Neanderthals had thick

phalanges, Ken thinks that they would have been unable to finger a primitive flute. Jen is not convinced. She considers the anatomical theory sketchy and its bearing on the case slight.

Many epistemologists follow William James in holding that our overarching epistemic objective is to believe as many truths as possible and to disbelieve as many falsehoods as possible (1983). They typically then focus exclusively on believing truths. But James's formula involves a proportion (see Riggs 2003). If an agent is epistemically risk averse, she will set a high threshold for acceptance. Perhaps she will accept only hypotheses that are .95 probable on the evidence. If she is more daring, she sets a lower threshold, perhaps .90. Both thresholds comply with the requirement that an acceptable conclusion must be highly probable. The risk averse agent will accept fewer truths and fewer falsehoods than the more risk tolerant agent. But neither is epistemically irresponsible. Jen is willing to accept a larger measure of epistemic risk than her mentor, Ben, is. She concludes that the artifact is a Neanderthal flute because, having taken all of the previous factors into account, she deems it 92% probable that it is. Ben, being more conservative, suspends judgment because he demands a probability of .95. He agrees with all of Jen's assessments. He simply does not think that they yield a high enough probability.

Evidently peers can differ along a variety of epistemically relevant axes: the weight to assign to evidence, the standards of acceptability, the identity of misleaders, the relevance and importance of various bits of background information, their favored styles of reasoning. To limit the cast of characters, I described Jen and Ken as differing along all these axes. Ben was introduced to show that even with widespread agreement in numerous other respects, epistemic agents can differ over thresholds of acceptability. But the factors I've mentioned vary independently and can point in different directions. Someone who sets a high threshold acceptability could easily favor or disfavor analogical reasoning. Along each axis, both Jen and

Ken hold reasonable views. Neither is making an obvious mistake. Nor is it plausible that either is making a subtle mistake. It is unreasonable to expect there to be a precise numerical value for how risk averse one should be; it is implausible to think that there is an algorithm for identifying misleaders, for assigning weights to different bits of evidence, or for assessing the credibility of an analogical argument.

We might think that this just shows that the epistemologists who set the criteria for peerhood should include more constraints. They should insist that epistemic peers also agree about the weight of evidence they assign, the criteria of acceptability, the reasoning strategies they actually use, and so forth. Then we would get back to the original question: should they be steadfast or conciliatory? The same arguments could be given on each side with a bit more precision. Once we control for all the differences I indicated (and probably others that I overlooked) , the familiar alternatives reappear.

The rationale for the original constraints on peerhood was to focus attention on disagreements that are epistemologically telling. As we have seen, not all are. Drawing on legal terminology, let us say that a factor is dispositive when it settles how a disagreement should be resolved. In some cases, imbalances in levels of expertise are presumptively dispositive. Once we recognize that one party is a novice and the other is an expert, or one party has a relevant cognitive asset and the other a corresponding deficit, it is usually clear how the disagreement ought to be resolved. But in disagreements between epistemically competent peers, there are no (even presumptively) dispositive disparities. To be sure, if we think the current criteria have omitted some dispositive differences, we can augment the criteria. Elsewhere I suggested, for example, that besides being equal in the respects standardly recognized, epistemic peers should be equally well educated in the relevant disciplines (2010). But the grounds for disagreement

that we have canvassed here seem different. In each of the respects we looked at, it is not at all obvious that either of the protagonists is in an epistemically superior position. We can determine how many additional false positives and false negatives will result from taking .90 rather than .95 as the threshold for acceptability, but that by itself does not tell us where the line should be drawn. Setting a precise threshold for how probable a conclusion should be, or how much weight should be attached to a given bit of evidence seems arbitrary. Moreover, it is apt to be counterproductive. Philip Kitcher (1990) argues that the scientific community may best serve its collective epistemic ends by, at any given time, countenancing a range of conflicting views. When there is a significant chance that a currently disfavored view is true, if the members of the community want to believe what is true, they ought not foreclose inquiry prematurely. Multiple ways of balancing the value of believing truths against the disvalue of believing falsehoods are often permissible. There is no reason to think there is a single optimal balance.

I suggest that we look at the availability of these sorts of disagreements as epistemic assets. The diversity of ways in which peers (as originally characterized) can still reasonably disagree should sensitize us to the complexity of a situation, and the epistemic opportunities that are available through taking different perspectives on it.

We have seen that judgments can be sensitive to choices among epistemically acceptable alternatives. Had an agent made a different choices within one or more acceptable ranges, she would have justifiably come to a different conclusion. This highlights the precariousness of our epistemic condition. That precariousness is independent of disagreement. The dependence of Jen's opinion on choices would be there even if she knew nothing about Ken's opinion -- indeed, even if Ken had no opinion on the matter. This might suggest that my discussion just underscores our fallibility. Despite the agent's best efforts, she still could be wrong.

Perhaps so. But disagreements among competent peers can be epistemically fruitful. If properly investigated, they provide resources for a focused fallibilism. If parties dissect their disagreement, they can discover how the various epistemically responsible choices affected their verdicts. Jen and Ben, for example, might come to appreciate that the basis for their disagreement has nothing to do with Neanderthals per se; it is about how great a risk of error paleoanthropology ought to tolerate. Given their mutual respect, disagreement might also provide parties with an incentive to rethink their choices. Ken might, for example, be prompted to reconsider whether he would view the matter differently if he revised his views about which elements of the evidence class are misleaders. Jen might rethink her doubts about the strength of analogical arguments. Even if all parties end up endorsing their previous positions, their stance vis à vis them is different. They appreciate that it is not just (or perhaps even mainly) a disagreement about the status of an artifact; it is also a disagreement about epistemic methods and standards that affect how one ought to determine the status of the artifact.

Open mindedness is a propensity to entertain alternative points of view. Not all points of view merit attention. None of our protagonists has any reason to think that the artifact in question is a space alien's cleverly disguised slide trombone. A competent peer's disagreement effectively certifies that a perspective is worth taking seriously. It raises the question: What is to be said for his way of looking at things? Given that he is a peer, the presumption is that there is something to be said for it.

The picture I have sketched about the value of disagreement has an additional payoff. It vindicates the way we do philosophy. Neither the steadfast nor the conciliatory really engage. The steadfast agent ignores the disagreement and holds fast to whatever she already believed. The conciliatory agent immediately backs off. (She may then recalculate, or recalibrate, but it is

open to her to simply suspend judgment and lower her credence without doing anything else.) Few if any philosophers exhibit either sort of behavior. If anyone were so bold as to simply dismiss a worthy opponent's disagreement out of hand, concluding that he must be wrong somewhere but there is no premium in attempting to figure out where, he would not be doing philosophy. Philosophy is not dogmatic. Someone who immediately suspended judgment on learning that a peer disagreed would not be doing philosophy either. Philosophy is not spineless. In fact, a lot of us are at home with disagreement, but rather uncomfortable with agreement. If someone says, 'I find your position entirely convincing', we are apt to think (if not say), 'Then you weren't really paying attention.' Faced with disagreement, philosophers engage. We attempt to uncover the basis of the disagreement, and assess its merits. We look for shared or unshared presuppositions, common or diverging conceptions, flaws in our or our opponent's arguments. Some of us may be sufficiently convinced of our views that we are quietly sure that our interlocutor has made a mistake somewhere, but we take it that we need to figure out where. Others may be more open-minded, both about whose mistake it is and even whether there is a mistake. In any case, we take it that it is our responsibility to take the disagreement seriously.

Moreover, we take it that there is considerable epistemic benefit in doing so. We learn more about our own commitments and their vulnerabilities. We learn more about alternatives to our position and the trade-offs that favor one side or another. Figuring out exactly where and why you disagree with David Lewis or Judith Jarvis Thomson is an intellectually rewarding experience. Taking well supported philosophical disagreements seriously is intriguing, informative, and fun. If we took our own practice as a model for how disagreement in general ought to be approached, we would be better off.

If we recognize that disagreement does not always indicate that someone has made a factual error, we can introduce a further stance -- one that I think philosophers take almost by default. Parties might remain *committed* to their own position, considering it a good basis for inference and action when their ends are cognitive (see Cohen 1992, Elgin 2017). Nevertheless, they recognize its vulnerability. This recognition is not idle. It requires them to be open to, and perhaps even to seek out, emerging counter-evidence and counter-arguments. This blocks Kripke's paradox. It gives them a strong incentive to strengthen and stabilize their position and to protect themselves should it turn out that they are wrong. It undermines complacency and fosters intellectual respect. Like the steadfast, the committed retain their convictions; unlike the steadfast they do not dismiss or disregard disagreement. Rather they use it, not only as a reminder that they might be wrong, but also as a probe, to tease out exactly where and why they might be wrong.

If the parties to the disagreement entertain one another's position seriously and respectfully -- if they keep a genuinely open mind to the possibility that there is something significant to be said for it -- they may come to appreciate weaknesses in their own position as well as strengths in their adversary's. This may lead one to concede that the other party is correct. Alternatively, it may lead her to shore up her own position so that it can deflect his objections. A third possibility is that together they craft a position different from the one either of them started with, or revise their methods, standards, or criteria of acceptance. There is no assurance that they will ever agree. Nor, as Kitcher (1990) argues, should we think that we would be better off if we could devise criteria that insured agreement in either the short or the long run. But by entertaining respectable alternative positions and appreciating their merits, we understand our epistemic situation better. Moreover, if we appreciate the basis of our

disagreement, we are in a better position to introduce appropriate safeguards. Given our fallibility, it is in general wise to take into account the possibility that a well supported opinion might still be false. We hedge our bets, take out insurance, make backups, diversify our portfolios, carry umbrellas in the recognition that our confident expectations do not always pan out. Attention to disagreement can pinpoint the places where safeguards are most called for. In cases like the ones we've discussed here, the safeguards in question might involve introducing caveats, or highlighting the ways a conclusion is based on assumptions that could legitimately be challenged. It might also involve efforts to stabilize one's position by showing that the conclusion is not too dangerously dependent on a controversial choice. If Jen could arrive at her conclusion without strongly depending on the disputed bits of flint, or without dismissing analogical arguments, her case would be stronger. If she sees how Ken's argument works against her, she can marshal resources for the defense. When we take disagreements among competent peers seriously, we gain a richer, more focused appreciation of our epistemic predicament. Whether we decide to remain committed or to conciliate, we are in a better position to appreciate the nature, scope and insights to be gleaned from our epistemic vulnerability.

References

- Christensen, David (2007). 'Epistemology of Disagreement: The Good News,' *Philosophical Review* 116:187-217.
- Cohen, L. Jonathan (1992). *An Essay on Belief and Acceptance*. Oxford: Clarendon Press.
- Elgin, Catherine (2017). *True Enough*. Cambridge: MIT Press.
- Elgin, Catherine (2003). *Persistent Disagreement* ed. Richard Feldman and Ted Warfield. Oxford: Oxford University Press, pp. 53-67.
- James, William (1983). *Principles of Psychology*. Cambridge: Harvard University Press.
- Kelly, Thomas (2005), 'Epistemological Puzzles about Disagreement,' *Oxford Studies in Epistemology* vol. 1, ed. Tamar Szabo Gendler and John Hawthorne. Oxford: Oxford University Press, 167-196.
- Kitcher, Philip (1990) 'The Division of Cognitive Labor,' *Journal of Philosophy* 87:5-22.
- Kripke, Saul (2011). 'On Two Paradoxes of Knowledge. *Philosophical Troubles*. Oxford: Oxford University Press, pp. 27-49.
- Riggs, Wayne (2003). 'Balancing our Epistemic Goals,' *Noûs* 37: 343-352.
- Sidgwick, Henry (1981). *The Methods of Ethics*. Indianapolis, Hackett.

1 Earlier versions of this paper were presented at the 2017 Voicing Dissent Workshop at the University of Connecticut, and at the 2017 meeting of the Canadian Philosophical Association. I am grateful to the participants in both for their help.

2 These equivalences are obvious idealizations. No two people are exactly alike in these respects. I follow the literature in making the idealizations, since the point is that differences of opinion are possible even when there is wide agreement in abilities and backgrounds.

3 I model this discussion on the controversy over the Divje Babe flute. I focus on whether it is Neanderthal, although there is also considerable controversy about whether it is a flute.