

The Evidentialist Theory of Disagreement

Brian Weatherson

1 Introduction

Philosophical debates about peer disagreement typically start with something like the following highly idealised, and highly stylised, case. Two agents, who we'll call Ankita and Bojan, have the same evidence concerning some question under discussion, and they know this. They each know that they are equally (and it's often implicitly assumed, highly) talented at resolving questions like this. Yet they disagree. For simplicity, we'll focus on the case where Ankita forms the doxastic states best supported in their situation, and Bojan does not. What should each do in the face of the disagreement?

Just as the debate has been structured around this idealised, stylised, case, so have the responses. Conciliationists say that each party should move their attitude, usually in the direction of the other's view. Anti-conciliationists deny this. The anti-conciliationist view I'm interested in says that Ankita is, in virtue of having got things right the first time, not obliged to move in the face of Bojan's error. This is the view that Thomas Kelly (2005) calls the Right Reasons view.

I'm not so much interested in defending the Right Reasons view as defending a more general view about disagreement that has Right Reasons as a trivial corollary. What I'm defending here is an evidentialist view of disagreement. Each party should change their views to the extent that the existence of the disagreement is evidence that there is evidence they have overlooked, as long as this second-order evidence (i.e., evidence about evidence) is itself evidentially relevant to the question under discussion. And they should, ideally, change their views to the view that is supported by the prior evidence, plus this new evidence of evidence. In the stylised, idealised case we started with, it is stipulated that each knows the other has the same evidence they do, so the disagreement provides no evidence whatsoever that there is new evidence, so they need not switch.

(Why does evidentialism support Right Reasons and not a Stand-Your-Ground view, where Ankita and Bojan should both stay still? Because it is hard to give univocal answers to what someone who has made a prior error should do. Assume that Bojan has looked at the evidence and concluded that JFK was murdered by Abraham Lincoln. What should he do in the face of new null evidence? I say that peer disagreement is null evidence, but we can consider something that's uncontroversially null evidence relative to this question, such as that it is sunny today. The best answer is equivocal. He should change his view, because he shouldn't think Lincoln murdered Kennedy. He should not change it on this basis, because the

sky's blueness is not a good reason to change his views on JFK's murder. I want to say the same thing about Ankita's disagreement with him; he should change, but not for this reason. I'll mostly side-step these issues by focussing on Ankita in what follows. Note that even some conciliationists, such as David Christensen (2011), are similarly equivocal about what Bojan should do in the face of disagreement.)

The evidentialist view of disagreement I'm promoting bears an obvious affinity to the justificationist view of disagreement that Jennifer Lackey (2010) defends. Indeed, I think the main differences are points of emphasis, not deep principle. Lackey describes her view as a way of taking the best features of each of conciliationism and anti-conciliationism; I'm interested in a version of the view that is clearly anti-conciliationist. Relatedly, Lackey's explanation of some of the cases that motivate conciliationism is different to mine. But the similarities outweigh the differences, and I wanted to note her theory as the closest precursor to the theory I'll defend here.

Another big motivation for the evidentialist view of disagreement I'm defending comes from some remarks on testimony by Frank Jackson (1987). Jackson suggests a broadly evidentialist treatment of testimony.

Why should you ever accept what I say, unless you already did so before I spoke – in which case speech is a luxury? ... The answer cannot be that you are taking me to be sincere. ... Sincerity relates to whether you should infer prior agreement or disagreement in beliefs, not to whether posterior adjustment of belief is in order. The reason posterior adjustment in belief may be in order is that hearers (readers) sometimes have justified opinions about the evidence that lies behind speakers' (writers') assertions. You assert that P. I know enough about you, and your recent situation, to know (i) that you have evidence for P, for you would not otherwise have said it, and (ii) that your evidence is such that had I had it, I would have believed P. I borrow your evidence, so to speak. Typically, I won't know exactly what your evidence is. Perhaps you visited a factory and came back and said 'The factory is well run'. I don't know just what experiences you had there – just what you saw, heard, smelt and so on – but I know enough to know that had I had these experiences – whatever exactly they were – I too would have come to believe the factory well run. So I do. ...in this way an epistemological division of labour is achieved. Imagine the work (and invasion of privacy) involved if we all had to duplicate each other's evidence. Of course, I may not come to believe exactly what the speaker or writer believes. A friend returning from overseas may say to me of a certain country 'It is very well run'. I may know enough of my friend to know

that experiences that would make him say that, are the kind that would make me say ‘Dissent is suppressed’. In this case, I will borrow his evidence to arrive, not at what he believes, but at what I would have, had I had his experiences. (Jackson, 1987, 92–3)

I agree with almost all of this, though I’m not going to defend such an evidentialist account of testimony here. (Why ‘almost’? Because it will be rather important later that we not be able to move as freely between sharing experiences and sharing evidence as Jackson does in the last line.) Rather, I’m just going to acknowledge my debt to Jackson’s ideas, and move to disagreement.

I’m hardly the first person to start with evidentialist intuitions and end up with anti-conciliationist conclusions about disagreement; you can see a similar trajectory in recent work by Maria Lasonen-Aarnio (2013, 2014), and what I say here also owes a lot to her. But I think the details are different enough to justify a new variant on similar themes.

2 Two Clarifications

To start with, note that my setup of the Ankita/Bojan case is rather ambiguous at one point, and philosophically loaded at another. The ambiguity comes in when I said that Ankita and Bojan are equally good at resolving questions like this. There are two natural ways to interpret this. We could read it as meaning that they are equally likely to come up with a rational verdict, or that their verdicts are equally reliable. David Christensen (2014) is very good on the importance of this distinction.

The literature typically concentrates on people one has (independent of one’s views on the disputed issue) good reason to take as *epistemic peers*—as rough equals along certain dimensions of epistemic evaluation. One such dimension concerns the evidence the other person has relevant to the disputed issue, and the other concerns how well she forms beliefs on the basis of her evidence. ...[W]e should notice that there are a couple of different ways of approaching the second dimension of evaluation—ways which are not always clearly separated. One focuses on the other person’s equal likelihood of responding *rationally* to her evidence. On this reading, ... the disagreeing friend is what might be called a “rationality-peer” on the given issue: one whose opinion is equally likely to be rational. The second way of evaluating the other person’s responses to evidence is in terms of her likelihood of responding to that evidence by forming *accurate* beliefs. On this reading, ... the disagreeing friend ... might be called an “accuracy-peer” on the given

issue: one whose opinion on the disputed issue one expects to be as likely to be accurate as one's own. (Christensen, 2014, 3)

Christensen cites Feldman (2007), Kelly (2005), Christensen (2007) and Cohen (2013) as writers who discuss peer disagreement using a rationality-based understanding of peerhood, and Elga (2007), White (2009), Enoch (2010), Kelly (2010), Lam (2011) and Levinstein (2013) as writers who focus on an accuracy-based understanding. And he's not the first to notice these two possible understandings; the distinction plays a big role in the arguments of Levinstein (2013) and Schoenfield (2014).

If peerhood is understood in rationality-based terms, then one motivation to conciliate in light of peer disagreement comes via a two-step process. The conciliationist says that Ankita should do two things in light of Bojan's disagreement. First, she should use that disagreement as evidence that her initial view is irrational, then second, she should that fact as grounds for revising that first-order credence. This is a broadly calibrationist motivation for conciliationism. And I think the second step in it is very implausible for reasons well laid out by Maria Lasonen-Aarnio (2014).

But many writers have noted in recent years that the first step of this sequence is dubious too. The most that Ankita gets from Bojan's disagreement is evidence that some other view to hers is rational. It is a fallacy to infer from that that her view is irrational, unless we have as a background assumption that there is only one rational response to any given evidence. So it seems that the argument for conciliationism requires the thesis Roger White (2005) calls Uniqueness: that there is a single rational response to evidence. Whether this seeming is really correct is actively debated: see Douven (2009), Kelly (2010), and Ballantyne and Coffman (2011, 2012) for interesting moves in the debate. I'm going to mostly not take a stance on this, since I think the arguments for conciliationism have other weaknesses.

It might seem that once Ankita views Bojan as an accuracy-peer, issues about higher-order evidence aren't relevant to determining whether she should conciliate in light of her disagreement. After all, in that case Ankita has two pieces of evidence; her own judgment and Bojan's. By hypothesis, each of them are equally accurate. So she should act as if she had two measuring devices, one which said that p was true, and the other that said it was false. And in that case one should have no settled view about p .

But Ankita doesn't just have two pieces of evidence; she also has the evidence that led to her initial judgment that p . We only get to describe the case in ways that make it seem symmetric if we somehow have a reason to set that initial evidence aside. This point is well made by Kelly (2010). And it really isn't clear why that should be set aside. (Weatherson (2013) describes one possible reason to set the public evidence aside, but also raises doubts about its effectiveness.) So it matters

how we understand peerhood, but either way there is no quick route to conciliationism.

I've said that Ankita and Bojan have the same evidence, as is common in setting up the puzzle. But although many writers say they are focussing on cases where the two parties have the same evidence, what they really provide are examples where the two parties have the same public evidence, but do not have the same private evidence, and so do not that they have the same evidence in total. Whether two parties have the same evidence is a philosophically loaded question, one that will become quite relevant to the arguments of this paper. To get some sense of why this might be relevant, consider the following example, which might at first glance seem to be a counterexample to conciliationism.

Stars I

Ankita and Bojan are wondering how many stars there are. They both have the concept of a prime number, but they aren't familiar with Euclid's proof of the infinity of primes. In fact, they both suspect, given the decreasing frequency of primes, that they run out eventually. In the course of their research into the stars, they run into the Delphic Oracle, who is known to always speak the truth. The Oracle says "There are as many stars as primes". Bojan takes this to be evidence that the proposition *There are infinitely many stars* is probably false. But while reflecting on the Oracle's words, Ankita comes up with a good proof that there are infinitely many primes, and concludes that this proposition is true.

This is a case where Ankita should not conciliate in light of her disagreement with Bojan. She has a proof that there are infinitely many primes and Bojan does not. So she should not change her views. If we take evidence to be constituted solely by happenings in the external world, then we might think that this is a case where Ankita and Bojan have the same evidence, form different credences, and yet should not, contra conciliationist dictates, conciliate. But that's not the right way to think of the case, and when we think of it the right way the case can easily be seen not to be a counterexample to conciliationism. Ankita has, by the time she concludes that there are infinitely many stars, more more evidence than Bojan. Her reconstruction of Euclid's proof is a bit of evidence she has, evidence which is highly relevant to the question under discussion, and which Bojan does not.

This all suggests a very weak, and hence easier to defend, version of conciliationism. It only applies to cases where two parties have differing views about a proposition, and the following four conditions are met.

1. The two parties have no reason external to this disagreement to think that one is more likely to be rational than the other.
2. The two parties have no reason external to this disagreement to think that one is more likely to be accurate than the other.
3. The two parties have the same public evidence.
4. The two parties have the same private evidence.

The sense in which I'm an anti-conciliationist is that even in this extreme case, I don't think both parties should conciliate. If one party is responding rationally and the other is not, the first party should stick to their view. But these conditions are met in vanishingly rare circumstances. And when they are not met, there are quite mundane reasons for thinking that each party should conciliate. A running theme through this paper will be that the cases thought to motivate conciliationism do not satisfy these four criteria, and hence it is possible for an anti-conciliationist to consistently say that each party should move towards the others view. So while the evidentialist theory of disagreement is anti-conciliationist in theory, it is somewhat conciliationist in practice.

Two more points of clarification before we move on.

First, I'm going to start by looking at a very specific form of conciliationism, namely Adam Elga's Equal Weight View (EWV). The EWV says that when two people are peers, and they have the same evidence, and they learn that they have credences c_1 and c_2 in a disputed proposition p , they should each adopt a credence half-way between their initial credences. That is, their new credence in p should be $\frac{c_1+c_2}{2}$. The EWV is not by any means the only version of conciliationism. Indeed, I think it faces some pressing technical problems, some of which are described by Jehle and Fitelson (2009) and others by Levinstein (2013). But it is instructive to separate the challenges for EWV from challenges for conciliationism in general.

Second, it is very important here, as almost everywhere in epistemology, to respect the distinction Gilbert Harman (1986) draws between inference and implication. We can see this by looking at a variation on the earlier example about stars and oracles.

Stars II

In this case Ankita and Bojan are very knowledgeable about primes. Indeed, they are among the co-authors of the counterpart paper to Polymath (2014). This time the oracle tells them that there are as many stars as twin primes. Ankita infers that there are probably infinitely many stars, but it is too soon to be completely confident. Bojan, on the other hand, becomes completely certain that there are infinitely many stars.

In my opinion, and for that matter Ankita's, the evidence Ankita and Bojan have conclusively settles the question of whether there are infinitely many stars. What they know about primes, plus what they know about the oracle, plus what they are told by the oracle, probably entails that there are infinitely many stars. So there is, probably, a conclusive implication from their evidence to that conclusion. But there is not by any means a reasonable inference from their evidence to the conclusion that there are infinitely many stars. That inference requires knowing something that is well and truly not in evidence, namely that there are infinitely many twin primes. The fact that this fact (and I think it probably is a fact) is a logical truth (or at least is logically entailed by things they know about primes) is irrelevant. A probably conclusive implication can be a definitely unreasonable inference, and is in this case. Unless Bojan has a proof of the twin prime conjecture up his sleeve, one that he hasn't shared with his co-authors, he should move his credences in the direction of Ankita's. That is, he should conciliate. It's possible that Ankita should conciliate too; I haven't said nearly enough about the case to settle that one way or the other. I think the mistaken idea that entailments generate maximally strong inferences has led to some confusion about what to say about certain cases, and that will become relevant as we progress.

3 Independence and Conciliationism

In early writings on conciliationism, such as Elga (2007) and Christensen (2009), there was a line of argument from Independence principles to conciliationism. This line is flawed, for reasons well set out by Errol Lord (2014). The point of this section is simply to rehearse Lord's arguments before moving onto other possible motivations for conciliationism.

There are weaker and stronger versions of the Independence principle. Independence principles in general say that judgments about the accuracy of one's judgment in a particular case can only be based on grounds independent of the original first-order judgment. When we think about disagreements, we're interested in second-order deliberations, hopefully ending in second-order judgments, that are prompted by disagreement at the level of first-order judgment. The strongest such principle says that in any dispute, a party to the dispute can only reasonably conclude that the other party is wrong based on reasons independent of their reasons for having a disputed view. But that leads to very odd predictions in cases like this one.

Bus Stop

While waiting at the bus stop, Ankita is approached by Bojan, who tells her that he is certain she lives in a shoe. Ankita is fairly confident, based

on long familiarity with her apartment, that she lives in an apartment, not a shoe.

Ankita doesn't have to find independent evidence that Bojan is mistaken to hold onto her belief that she lives in an apartment. Perhaps in some realistic versions of *Bus Stop*, Bojan would appear drunk or be slurring his words, and that would be the relevant independent evidence. But those external clues are not necessary. Bojan could appear perfectly sane and sensible in every respect except his firm belief that Ankita lives in a shoe, and she could still dismiss his view. So this strongest independence principle is false.

More plausible independence principles restrict the circumstances in which one must rely on independent reasons to dismiss a conflicting view. There are two interesting restrictions we could look at:

- Independence might be restricted to cases where the disagreeing parties are known to be just as good at reading the evidence. (We could break this down into two sub-cases depending on whether 'good' is understood in terms of accuracy or rationality, but this won't matter.)
- Independence might be restricted to cases where the disagreeing parties are known to have the same evidence.

But, and this I think is really the crucial point that Lord makes, neither of these restrictions on their own gives us a plausible principle. If we only impose the first restriction, we end up with the implausible conclusion that Ankita is expected to conciliate in the following case.

Party

Ankita and Bojan are just as good, in both senses, at working out where a party is given some evidence. But Bojan hasn't looked at the invitation to tonight's party in weeks, so is uncertain whether the party is on State St or Main St. Ankita looked at the invitation two minutes ago, and is certain the party is on State St.

It would be absurd to think that because Ankita's credence that the party is on State St is 1, and Bojan's is 0.5, and they are just as good at working out where parties are given some evidence, that Ankita's credence that the party is on State St should move to 0.75. Rather, she should conclude that Bojan hasn't looked at the invitation recently. And she should conclude that simply because Bojan has a different credence to her about where the party is. That's what an independence principle that

only imposes the first constraint would rule out, so such an independence principle is false.

Nor will the second restriction on its own do. If we restrict the restriction to public evidence, then Stars I is already a counterexample to it. But we can come up with cases where arguably Ankita and Bojan even have the same private evidence, and the restriction is still not sufficient.

Diagnosis

Ankita is a professor at a medical school, and Bojan a student. The students at her school are very good; often they are as good at diagnosis as the professors. And Bojan has done, Ankita knows, very well on his theory exams. But some students who know a lot of theory are very poor at making a diagnosis based on material in a patient's file. So Ankita pulls out a file at random for her and Bojan to look at it. Given the symptoms displayed, Ankita is very confident in a particular diagnosis. But Bojan has no idea what to say about the case; his best guess is that we should have low but positive credence in several distinct diagnoses.

In this case, Ankita shouldn't infer that she had been over-confident. She should conclude that, despite his solid background, Bojan isn't very good at making a diagnosis. I've obviously simplified a lot, but this seems like a very natural way for professors to test that their students have or lack a practical skill. Now perhaps the best diagnosis is that Bojan really lacks some evidence, despite his doing well on tests. That's actually what I suspect is going on. But I suspect most people, and certainly most conciliationists, don't think that. It is much easier to motivate conciliationism if we think that there is a skill of processing evidence that goes well beyond the possession of evidence, and that in cases like this one what's happened is that Bojan lacks that skill. (Why say conciliationism is easier to motivate if one posits large skill differences that go beyond evidence possession? Because now we can say why one person should defer to another without thinking the other person has evidence they lack; the other person may have more skills.)

Now if independence just requires that the parties had the same evidence, and this is a case where the parties have the same evidence, it would be an independence violation for Ankita to infer from Bojan's lack of certainty in any diagnosis to his lack of skill in making diagnoses. Rather, she should conciliate with him, and lose confidence in her diagnosis. That's wrong, so this independence principle is too strong.

So just putting each of these restrictions on independence singularly does not yield a viable principle. What happens if we put both restrictions on at once, and say independence holds only if the parties are known to have the same evidence and

known to be just as good (in some sense) at processing it? Lord points out that then we don't have a premise in an interesting argument for conciliationism. Rather, the independence principle that is supposed to motivate conciliationism has just become a statement of conciliationism. So it can't provide any independent support for it.

4 Six Examples

For these reasons, I accept Lord's conclusion that it is not possible to provide a good argument based on independence principles in favor of the EWV, or indeed of any conciliatory position. But I also don't think that these principles have been what have most moved people towards these conciliatory positions. Rather, they are moved by the idea that the EWV, or at least some form of conciliationism, is the best explanation of the clear facts about some simple cases. I think that the literature here, as with the literature on higher-order evidence, suffers from that "main cause of philosophical disease—an unbalanced diet: one nourishes one's thinking with only one kind of example." (Wittgenstein, 1953, §593). I can't claim to offer a balanced diet, but I can offer the start of a more varied one. Here are six new morsels that will form the basis of the discussion to follow. I'm going to argue that the evidentialist can explain the last two better than any conciliationist can, and there is no case that the conciliationist can explain better than the evidentialist can. So to the extent that our views about disagreement should reflect how well the theories can handle cases like these, we should be evidentialists. I think that considerations other than reflection on cases also favour evidentialism over any kind of conciliationism, but I'm not going to argue for that here. Instead, I'll launch into the cases.

4.1 Arithmetic

Ankita and Bojan are working on some arithmetic problems. They both know that they have a similar track record at these problems; both are reliable, with very similar rates of mistakes. They are trying to work out *What is 22 times 18?* Ankita correctly works out that it is 396; Bojan says that it is 386. What should their credences in each answer be?

4.2 Jellybeans

Ankita and Bojan are trying to guess how many jellybeans are in a sealed, transparent container. They both have equal access to the container, and they both know that they have similarly good track records at this kind of game. Ankita correctly guesses that there are 396; Bojan guesses that there are 386. What should their credences in each answer be? (A similar case is considered by Jack L. Treynor (1987).)

4.3 Detectives

Ankita and Bojan are the two best murder detectives in the world. They both know that they are the only peers they each have, and that they have very similar track records of success, with equal (and rare) failures. They are brought in to solve a mystery that no one has made any progress on. Each quickly sees that it could only be the butler or the gardener. Bojan has equal credence in each suspect, but Ankita figures out a subtle reason that it could not have been the gardener, so is sure the butler did it. And in fact the butler did do it, and Ankita is right about why the gardener could not have done it. After they compare credences, Bojan giving equal credence to each suspect, and Ankita being sure it is the butler, what should their credences in each answer be? (I owe this case to Ben Levinstein (2013).)

4.4 Football

Ankita and Bojan are both very good at predicting football games of different codes, and they both know this, and they have the same public evidence. They are comparing their credences in the home team winning ahead of two big matches: an Australian Rules match in Melbourne, and an English Premier League match in London. For each match, Ankita has a credence of 0.9 that the home team will win, and Bojan has a credence of 0.1 that the home team will win. They both regard the matches as completely independent, so Ankita's credence that both home teams will win is 0.81, while Bojan's is 0.01, and each of them have credence 0.09 in each of the hypotheses that one particular home team will win and the other will not. Once they share their credences with each other, what should their credence be that (a) the home team will win the Australian Rules match, (b) the home team will win the English Premier League match, and (c) both home teams will win?

4.5 Simple Arithmetic

Ankita and Bojan are working on some arithmetic problems. They both know that they have a similar track record at these problems; both are reliable, with very similar rates of mistakes. They are answering the question *What is 2 plus 2?*. Ankita says it is 4; Bojan says that it is 5. What should their credences in each answer be?

4.6 Doctors

Ankita and Bojan are the two best cardiologists in the world. They know each other to be peers, the only peers each has. They are brought in to diagnose a case that has stumped all the other experts in the field. Ankita judges that it is likely disease A, but she is just short of fully believing it is disease A, since she thinks disease B is an unlikely, but real, possibility. This is the rational response to the evidence. Although the patient has disease A, the public evidence is just short of being sufficient to ground knowledge that the patient has disease A, since B is also

a realistic possibility. She reports all this when she and Bojan compare notes, but Bojan reports that he is confident that the patient has disease A. What should their credence in each diagnosis be?

4.7 My Verdicts

These cases are, in general, not so clear that we can simply know what is true about them after a moment's thought, and use that knowledge to evaluate theories. But for the record, here are my verdicts on the cases.

In **Arithmetic**, I think a lot depends on the finer details of the case, particularly on how Ankita got to her answer. But I think on the most plausible filling in of the details, there isn't a lot of pressure on her to conciliate. Now this isn't a popular view. Much of the motivation for conciliationism comes from thinking that in versions of **Arithmetic** where the sum in question is not specified, there is rather strong pressure to conciliate. But as we'll see, we really need to look at the quite fine detail of cases like **Arithmetic** to say what Ankita should do.

In **Jellybeans**, I think they clearly should conciliate. And unlike in **Arithmetic**, this conciliation should take the form of not just lowering their credences in their preferred answers, but in increasing their credence in answers between the two they offered. In **Jellybeans**, the announced answers should increase their confidence that the answer is 391, which is not what should happen in **Arithmetic**.

I have no idea what the answer to the third question in **Football**, about the appropriate credence in the compound proposition, is. We'll say a bit below about why this is such a hard question.

In each of the last three cases, **Detectives**, **Simple Arithmetic** and **Doctors** Ankita should not conciliate, and Bojan should move his credence dramatically in the direction of Ankita's. Or at least so I say.

5 Equal Weight and the Cases

On the face of it, the EWV gets at most one of the six cases right. After all, the only case where it seems even *prima facie* right to move to a credence half-way between the two expressed views is **Arithmetic**. But I think matters are not so simple. A more nuanced understanding of the cases lets EWV handle **Jellybeans**, and a more subtle version of conciliationism does well (or at least well enough) with **Detectives** and **Football**. If there is a case-based objection to conciliationism, it comes from the last two cases. But first I want to go over why the second, third and fourth cases are really not problems for conciliationism. Why, as an anti-conciliationist, should I do that? It's for two reasons. First, I want to demonstrate how hard it is to use any case around here to show that a particular view on disagreement is wrong. Second,

we get an interesting insight into the range of possible and indeed plausible versions of conciliationism by working through the cases carefully.

The apparent problem with **Jellybeans** is that it seems the rational reaction, for both Ankita and Bojan, is to increase their credence in a particular hypothesis that neither of them endorses, namely that there are 391 beans in the jar. But this is a merely apparent problem. What credences is it rational to attribute to Ankita when she announces her guess of 396? Presumably it isn't that she has credence 1 that there are 396, and credence 0 in everything else. Given what we know about jars of jellybeans, and human visual capacities, it is best to interpret her as saying that the mode of her credal distribution over the competing hypotheses about the content of the jar is 396. But that distribution will presumably be fairly spread out, and indeed fairly flat around the peak. Similarly, Bojan will have a credal distribution that is spread out, and fairly flat, around its peak at 386. If we average out those distributions, it could easily be that the peak of the new distribution is at 391. That happens, for instance, if each of Ankita and Bojan's distributions are normal distribution, with a mean at the number they announce, and a standard deviation of 10.

Now there are hard questions about how we do, or even could, know that the number they utter means that they have just this credal distribution. But that's not particularly our problem here. The question is what the parties to the dispute should do given that we add to their evidence each other's credence distribution. Questions about how we could know what another person's credence distribution are, while fascinating, are not at issue here. This is a point worth keeping in mind as we work through the examples.

The EWV also does rather badly on **Detectives**. Assuming that the detectives are actually very good at their jobs, then neither would have formed the conclusion that the butler did it without a very good reason. If one of them believes this, and the other does not, the one who does not should believe that they've missed a reason. So they should largely defer to the other.

Note that the reasoning in the last paragraph is entirely symmetric, and doesn't directly make use of the fact that Ankita was right to infer that it was the butler. So it is reasoning that should be available to the conciliationist, even if it isn't available to the equal weight theorist. And I think there is a natural method for how to make sense of **Detectives** in a conciliationist-friendly way. The method in question is one I'm taking from some work by Sarah Moss (2011).

Imagine that Chika is not a detective, and has no particular expertise in solving murders. Moreover, she has very little information that bears directly on the case. She does have one very important skill; she is great at judgment aggregation. And while that's a rare skill to have, it's exactly what is needed here, since she has no

relevant evidence about the murder apart from knowing what Ankita and Bojan think. That is, she knows that Ankita is confident the butler did it, while Bojan is uncertain. Now the reasoning from two paragraphs ago is available to Chika too. She can think that Ankita wouldn't be so confident unless she had a very good reason, so she can infer that it is very likely that the butler did it.

One natural form of conciliationism says that the parties to a dispute face the same normative pressures as an outsider, like Chika. Whatever is rational for Chika to do given the knowledge just of the parties' credences, and their track records and backgrounds, is rational for the parties to the dispute to do. In general, that will mean conciliating, since in general Chika should form a credence somewhere between the parties' credences. But perhaps that isn't always true. If Ankita and Bojan were both 90% confident that it was the butler, and that's all Chika knows, then Chika should give some credence to the possibility that Ankita and Bojan have noticed independent reasons for thinking it is likely to be the butler, and so should have a credence in the butler's guilt slightly higher than 0.9. Nevertheless, the view that insiders to the dispute, like Ankita and Bojan, should end up in the same place as an outsider, like Chika, who knows just the credences, seems to capture the idea at the heart of conciliationism.

I haven't said very much in general about how Chika should reason about these cases. Ben Levinstein (2013), to whom I owe this example, thinks that Chika should reason in a way that is sensitive to what Ankita and Bojan's expected inaccuracies will be given her subsequent credences. Levinstein's arguments strike me as very persuasive, and I encourage anyone more interested in the question of how Chika should solve problems like this to study Levinstein's work.

We can apply the same approach to **Football**, which is actually a rather hard problem. I used to think that it was a counterexample to any form of conciliationism. Thinking about conciliationism as an 'insider-outsider' thesis, as just described, made me realise that isn't true. What's true about **Football** is that it is hard; hard for Ankita and Bojan, but hard for Chika too. Indeed, I have no idea what she would, given her assumed skills, think in this situation. Let's look at why it is so hard.

Assume, more or less for reductio, that Chika, our ideal aggregator, aggregates probabilistic judgments by taking the linear average of them. (If that's right, the EWV and the 'insider-outsider' version of conciliationism coincide; if it isn't right, they don't.) The following table gives Ankita, Bojan and Chika's credences and conditional credences, assuming that Chika does this. I'll use p for the home team wins the Australian match and q for the home team wins the English match.

	p	q	$p \wedge q$	$q p$
Ankita	0.9	0.9	0.81	0.9
Bojan	0.1	0.1	0.01	0.1
Chika	0.5	0.5	0.42	0.84

The key number is in the bottom right. Assuming that Chika plans to update by conditionalisation, that means that although her credence in q is now 0.5, if she learns p , it will rise to 0.84.

It has been argued, e.g. by Loewer and Laddaga (1985) and Jehle and Fitelson (2009), that this is a mistake for the following reason. Ankita and Bojan both take the games to be probabilistically independent. So Chika, who only has their credences to go on, should take them to be independent too. This argument doesn't work, for a reason Sarah Moss (2011) gives. The probabilities in this table are evidential probabilities. Even if the games are physically independent, it could be that the result of one gives Chika evidence about the other. And that is what happens; if she learns p she gets one more data point in favor of Ankita's general accuracy in football-predicting, and against Bojan's. So it is plausible that, for her, learning p will raise her credence in q .

What isn't plausible, however, is that it could raise her credence that much. We can imagine, consistent with everything I've said so far, that Chika has a lot of evidence about Ankita and Bojan. One result going the way Ankita predicts might break a tie between them in terms of long-run accuracy, but given that kind of background, it can hardly make it rational for her to move almost all the way to Ankita's view. If she learns p , she'll be in a position where she has advice about q from two experts, one of which has an ever so slightly better track record. It is implausible to think that she should let this mild asymmetry move her so closely to that expert's view. More generally, it is implausible that the value in the bottom right should be completely independent of just how much evidence Chika has about Ankita and Bojan's track records. Arguably, if she is very poorly informed about each of them, then learning that Ankita got one game right and Bojan got it wrong could have this kind of dramatic difference. What is less plausible is that how much she should move her credence in q in light of this new evidence about Ankita's accuracy is independent of how much evidence she had about Ankita (and Bojan) prior to the two games being played. Yet that's what the Equal Weight View says.

So the Equal Weight View is simply wrong about this case. But conciliationism need not be. Saying just what values should go in the two right-most boxes in the bottom row is a very very hard question. But it is implausible that the question has no good answer; that somehow Chika is thrown into something like an epistemic dilemma here. And the conciliationist who says that whatever is right for Chika to

do, that's right for Ankita and Bojan to do once they learn of each other's views, is not saying anything incoherent or obviously implausible. I don't think this conciliationist view is ultimately correct, but I wouldn't want to rely on intuitions about the case to say so. Saying that Ankita and Bojan face the same kind of hard judgment aggregation problem as Chika faces seems, *prima facie*, relatively plausible.

But the same move isn't as plausible for **Doctors**. If Chika knows that Bojan is certain of a diagnosis, and that Ankita gives that credence a very high credence just short of belief, it seems *prima facie* plausible that Chika should conclude from that that the diagnosis is correct. Unless we have some way to motivate a theory of judgment aggregation where the aggregate opinion is never more confident in a proposition than the weakest member, there must be some such cases where Chika should believe the diagnosis is correct. But Ankita should not share this confidence. She should not find her doubts assuaged by Bojan's not sharing them.

It is easy for the evidentialist to say what's going on in **Simple Arithmetic**. Ankita has maximally strong evidence that 2 plus 2 is in fact 4. That's not just because the conclusion is a logical truth. There are plenty of logical truths that we have insufficient evidence to believe.¹ Rather, it is because the inference from $x = 2 + 2$ to $x = 4$ is one that is immediately justified, without the need for further steps. Bojan's disagreement can't dislodge that.

But how can the conciliationist handle the case? It doesn't seem very plausible to say that when an otherwise reasonable person says that two plus two is five, we're obliged to doubt that it is four. The usual response on behalf of conciliationists is to appeal to the notion of 'personal information'. The idea was first developed by Jennifer Lackey (2010), but I want to first mention the version of this defence put forward by David Christensen (2011). (Christensen is describing a scenario where the narrator plays the role of Ankita, and Bojan is their friend.)

If such a bizarre situation were actually to occur, I think one would reasonably take it as extremely unlikely that one's friend (a) was feeling as clear-headed as oneself; (b) had no memories of recent drug-ingestions or psychotic episodes; and most importantly, (c) was being completely sincere. Thus, to use Lackey's term, one's personal information (that one was feeling clear, lacked memories suggesting mental malfunction,

¹Do we have sufficient reason to believe every instance of excluded middle? Let me get back to you on that after I've carefully studied a few hundred papers on inferentialism about logical connectives, the importance of harmony, the structure of logical consequence relations, the best characterisation of the realism/anti-realism debate, and other related topics. Maybe it is true these are immediate logical implications of the null evidence set. But maybe they aren't; it's too soon to tell. I'm confident that I couldn't know I'm in a position to know every instance of excluded middle; but maybe I really do have compelling evidence that each is true. These questions are hard.

and was being sincere in one's assertion) would introduce a relevant asymmetry, and one could reasonably maintain one's belief.

The first thing to be said here is that (c), which is what Christensen adds to Lackey's original characterisation, is beside the point. The question is what Ankita should do given that Bojan believes that 2 plus 2 is 5. It's not the separate question of whether she should believe he believes that, given his utterance. So questions of sincerity are beside the point. Then the question is whether (a) and (b), which are the aspects of personal information that Lackey originally highlighted, are enough to help.

And it is hard to see how they could be. If the reason for discounting Bojan's opinion rested on one's personal information, then the more information we get about Bojan, the more worried we should be. But I rather doubt that running a drug test on Bojan, to see whether (b) is a relevant difference between him and Ankita, should make any difference at all to Ankita's confidence.

More generally, this explanation seems to rest on an odd view about epistemic capacities. Ankita's ability to do simple arithmetic is not a sufficient ground to believe that two plus two is four. But her ability to detect differences in capacities and aptitudes between two people, one of whom is herself, is enough of a ground. Speaking personally, I'm sure I'm much better at simple arithmetic than I am at doing such a comparison between people. Indeed, my abilities to make such comparisons intuitively are so weak that I could only possibly do them by careful statistical analysis, and that would require, among other things, being able to add two plus two. In other words, if I can't know what two and two is, I can't process the evidence that might tell for or against the abilities of one party or another. So the conciliationist doesn't have a good explanation of how we can hold on to knowledge in these simple cases.

Simple arithmetic cases are important not just because they raise problems for conciliationism, but because they tell us something about what's at issue in debates about disagreement. Consider this argument by David Enoch for thinking that in debates about disagreement as such, we should treat the parties to the disagreement symmetrically.

Second, our question, as you will recall, was the focused one about the epistemic significance of the disagreement itself. The question was not that of the overall epistemic evaluation of the beliefs of the disagreeing peers. Kelly is right, of course, that in terms of overall epistemic evaluation (and barring epistemic permissiveness) no symmetry holds. But from this it does not follow that the significance of the disagreement itself is likewise asymmetrical. Indeed, it is here that the symmetry is so compelling. The disagreement itself, after all, plays a role similar to

that of an omniscient referee who tells two thinkers ‘one of you is mistaken with regard to p ’. It is very hard to believe that the epistemically responsible way to respond to such a referee differs between the two parties. And so it is very hard to believe that the epistemic significance of the disagreement itself is asymmetrical in anything like the way Kelly suggests. (Enoch, 2010, 657)

Well, consider the case when p is the proposition that two plus two is four, and Ankita is the party who believes p , while Bojan rejects it. Having an omniscient referee tell the parties that one of them is mistaken should produce asymmetric responses in the two parties. Now maybe there are only a small class of cases where this is the case, and what Enoch says is right in the majority of cases. But we can’t argue for that on perfectly general grounds about the nature of disagreement, because it fails in extreme cases. The argument that it holds in normal cases needs a distinct defence.

6 The Evidentialist Approach

Having gone over how conciliationism handles, or doesn’t handle, the cases, let’s compare it to evidentialism. We’ll look at them in reverse order, because the later cases are much easier for this evidentialist approach to disagreement.

The case gets **Simple Arithmetic** clearly right. Ankita has clear and compelling evidence that two plus two is four. In fact, the fact that two plus two is four is part of her evidence, and that’s as straightforward as evidence gets. Learning that something has gone badly wrong with Bojan’s arithmetic is hardly a reason to give up on this evidence.

The evidentialist view also gets **Doctors** right, though saying why requires a detour into views about higher-order evidence. Let’s start by considering the puzzle generated by these two cases.

Resident I (Christensen, 2010, 186)

I’m a medical resident who diagnoses patients and prescribes appropriate treatment. After diagnosing a particular patient’s condition and prescribing certain medications, I’m informed by a nurse that I’ve been awake for 36 hours. Knowing what I do about people’s propensities to make cognitive errors when sleep-deprived (or perhaps even knowing my own poor diagnostic track-record under such circumstances), I reduce my confidence in my diagnosis and prescription, pending a careful recheck of my thinking.

Resident II

Roshni is a hospital resident, with a patient in extreme pain. She is fairly confident that the patient has disease X, but thinks an alternative diagnosis of Y is also plausible. This is a rational response to the evidence; X is the most likely diagnosis, but Y is a realistic alternative given the known facts. The treatment for X would relieve the pain quickly, but would be disastrous if the patient actually has Y. Her judgment is that, although this will involve more suffering for the patient, they should run one more test to rule out Y before starting treatment. Roshni is then told that she has been on duty for 14 hours, and a recent study showed that residents on duty for between 12 and 16 hours are quite systematically too cautious in their diagnoses.

The narrator in **Resident I** (who I'll just call Narrator) should recheck his calculations. It would be wrong to act on unchecked judgment made by someone awake for 36 hours. But Roshni should not change her plan of action because of the higher-order evidence. It would also be wrong for her to simply treat the patient for X, just because most people in her situation are a little cautious in making a diagnosis. Those facts about what should be done seem pretty clear-cut to me. It's far from obvious, but I think they are best explained by facts about confidence. Narrator has a reason to reduce his confidence in his initial diagnosis, but Roshni does not have a reason to increase her confidence in the diagnosis of X.² We need an explanation of this asymmetry.

The explanation that Christensen offers of **Resident I** doesn't work. He says that in these cases, the residents should 'bracket' their first-order evidence, and just use the facts of their first-order judgments, plus facts about the circumstances in which those judgments are made. That would imply that Roshni should get more confident in the diagnosis of X, and that's not true. The evidentialist, by contrast, can explain each of the judgments.

In **Resident I**, if the case is at all realistic, it will involve abductive reasoning on the part of Narrator. And abductive reasoning is non-monotonic; it could be overturned by more evidence. So to make the diagnosis, Narrator needs to be in a position to reasonably believe claims like *There are no other facts that I need to learn before making a diagnosis*, and *There are no other plausible explanations of the data*. Now if

²Strictly speaking, it doesn't matter for the argument of this paper whether the claim I've just made is true. All I'm going to be using from the example is that the evidentialist can explain the asymmetry in how Narrator and Roshni should act. But I do think that there is an asymmetry in what they should believe, and evidentialism can explain it, and that's evidence in favour of evidentialism. That's why I'm talking both about the asymmetries in action and the asymmetries in belief between Narrator and Roshni.

the higher-order evidence about people who have been on duty that long is accurate, then Narrator's beliefs in these further facts is highly unreliable. And so they weren't part of his evidence. So he didn't have sufficient evidence to make the diagnosis. In the natural telling of **Resident I**, all that will be true before Narrator is reminded of how long he has been awake; he'll simply never have a justified belief. If the higher-order evidence is misleading, then Narrator will be forming these extra beliefs by a reliable method, but it will be a method he has good reason to believe is unreliable. And one can't get knowledge by using a method one has good reason to believe is unreliable. So Narrator still doesn't know the relevant background facts, so still doesn't have enough evidence to make the diagnosis. Note that in this case, Narrator may have originally known that the diagnosis was correct, but then lost the knowledge once he got reason to believe that his methods for forming beliefs in relevant background facts was unreliable.

Contrast this case with Roshni's. In order to form a belief that her patient has disease X, Roshni would need to have, as part of her evidence, that there are no reasonable alternative explanations of the data about the patient. And that means she needs to know that is true. But she can't know it is true, because it simply isn't true. It might be that she could be misled into reasonably believing that there are no reasonable alternative explanations, and that she is being unreasonable in being cautious. But in order to have sufficient evidence to believe the patient has disease X, she would need to have in her evidence something that is not true. And while higher-order evidence that one is excessively cautious can make a falsehood reasonable to believe, it can't make that falsehood into a piece knowledge. And without making it into knowledge, it can't make it into evidence. And that means Roshni can't get the evidence she needs to diagnose the patient with X.

Returning to disagreement, the evidentialist can explain what's going on with **Doctors** in the same way. To reasonably make the diagnosis of disease A, Ankita needs to know that there is not a viable alternative explanation. If she doesn't know that, she doesn't have enough evidence to make the abductive judgment. She doesn't know that before talking to Bojan, and talking to him doesn't make it true that there is no viable alternative explanation. So she shouldn't change her beliefs, or her actions. Note that this story is perfectly consistent with saying, in ordinary cases of disagreement between alternative diagnoses, that both doctors should conciliate before acting. It just says that conciliation with someone who is unreasonably more confident than you in a key proposition is unreasonable. And that's the fact to be explained in **Doctors**.

For the evidentialist, **Detectives** is a lot like **Doctors**. When Ankita hears Bojan's credence, it is reasonable for her to infer that she has some evidence that Bojan lacks. This evidence need not be public evidence; it might be more like the

kind of evidence a mathematician gets when working through a proof. But it is reasonable for her to infer, given just the facts about their conflicting credences, that Bojan has simply missed the reason that it must have been the butler. So she doesn't have new evidence that it wasn't the butler, so her credence shouldn't move.

It's a little harder to know what it should say about **Football**. The case as presented didn't include much detail about how Ankita or Bojan came to their conclusions. If I was in one or other of their positions, I would likely infer that the other had picked up on some reason I missed, but also that they had probably missed some reason I'd seen. So I would be tempted to conciliate, because this is a case where the conflicting credences really are useful evidence that there is (private) evidence that would motivate a change of view.

Note one consequence of this view that is, I think, entirely right. Ankita should move her credences about the games in the direction of Bojan's on hearing his view. But it would also make sense for her to ask Bojan why he has the view he does. And, if it turns out that all the reasons Bojan can offer are ones that Ankita had already properly weighed, she should revert to her initial credence. That all sounds like good practice to me. The primary reaction to hearing that someone has a very different view to one's own shouldn't be to jump to a new credence, it should be to find out why they have the conflicting view.

In **Football** it was plausible that the two parties would have different evidence; in **Jellybeans** it is just about certain. Ankita and Bojan will have had different appearances when they looked at the jar, they will have seen it from different angles, they will be bringing different histories with these kinds of estimation tasks to bear on the subject, and so on. In **Football** it was likely that the parties will have different views about the question because they have different evidence; in **Jellybeans** it is practically certain. So the evidentialist view says, along with intuition, that this is a case where they should conciliate. It has a simpler explanation as to why their credence in hypotheses like 391 should increase than the conciliationist offers, but both parties get to what seems like the right result for plausible reasons.

The big case that's left is **Arithmetic**. This case seems to be the one that moves people to reject evidentialist views. For instance, Alex Worsnip presents a version of **Arithmetic**, and announces on the basis of intuitions about it that views like evidentialism are "not even slightly plausible." (Worsnip, 2014, 6).³ So let's discuss it in more detail.

The first thing to note about the case is that it leaves out some details about how Ankita came to her conclusion, and filling in those details could change the case a

³In the quote in question, Worsnip is actually talking about what he calls 'extreme steadfastness'. But as he makes clear on, for example, page 16 of his paper, he is treating views like the evidentialist one I'm describing as versions of the steadfast view.

little. So let's present two variants of the case. The variations will be important enough that I'll introduce new characters to participate in them.

Deanna thinks to herself that 22 times 18 is 20 times 18 plus 2 times 18, so it is 360 plus 36, so it is 396. That strikes her as conclusive, so she announces that it is 396. Bojan then says he thinks 22 times 18 is 386. So Deanna decides to double check. She thinks that 22 times 18 is 20 plus 2 times 20 minus 2, so it is 20 squared minus 2 squared, so it is 400 minus 4, so it is 396. She now feels confident sticking to her original verdict.

Efrosyni thinks to herself that 22 times 18 is 20 times 18 plus 2 times 18, so it is 360 plus 36, so it is 396. But she has a little doubt about whether that reasoning was correct, and feels she should double check. So she thinks that 22 times 18 is 20 plus 2 times 20 minus 2, so it is 20 squared minus 2 squared, so it is 400 minus 4, so it is 396. She now feels confident sticking to her original verdict. She then hears Bojan say that he thinks 22 times 18 is 386.

It seems to me that, whatever one's view about how confident Deanna and Efrosyni should be in their verdict that 22 times 18 is 396, they should be equally confident. They both have two calculations they have done that point to that answer, and they have the same defeater, namely Bojan's conflicting judgment. But I don't see how any form of conciliationism can deliver that result. After all, conciliationism requires that a form of independence hold.⁴ The reasoning that led to one's disagreeing views cannot be used to 're-check' that those views are correct. So once Efrosyni hears Bojan's disagreement, she can't rely on either of the two routes to the conclusion that she used. But Deanna is free to use the second calculation she did as independent evidence that Bojan is wrong. So the standard conciliationist has to say, falsely, that Deanna and Efrosyni should have different credences in the proposition that 22 times 18 is 396, or, equally falsely, that Deanna doesn't get any reason to believe that 22 times 18 is 396 when she does the double-check.

The evidentialist has a better analysis of the case. Consider the state of mind that Efrosyni was in when she thought, "I'd better double check this." She actually had conclusive evidence that 22 times 18 was 396. Of course, she did that when she was making coffee and not thinking about arithmetic, so perhaps that isn't so important. What is more important is that after doing the first calculation she had evidence that a reasonable person could, other things being equal, base a belief on. Deanna was not unreasonable when she made her announcement, but yet Efrosyni in a similar position thought she should get more information. How should we explain

⁴The discussion of Lord's work above wasn't meant to undermine that claim. The result of that discussion was that conciliationism is equivalent to the strongest plausible independence principle, so that principle can't be used to independently defend conciliationism. That's all consistent with saying that conciliationism requires an independence principle.

that? We could treat this as a case where a kind of permissivism is right; Deanna was being reasonable in ending inquiry, and Efrosyni was reasonable in not ending it, despite their being in identical positions. But I think it is better to treat these cases as not quite identical. Efrosyni had a nagging doubt, which Deanna did not have. Perhaps that is the difference; the calculations they had both done are sufficient to end inquiry in the absence of positive reasons to extend inquiry. A nagging doubt like Efrosyni had is reasonable, and should be addressed. But it is also reasonable to not have such a doubt.

If that story is right, then the evidentialist can say what's going on in the disagreement case. Or, more accurately, in the disagreement cases. If Ankita is like Deanna, then the exchange with Bojan provides a good reason to recheck her calculations. The idea here is that the evidence is good enough to close inquiry, in the absence of positive reason to keep the inquiry going. That reason could be internal, a nagging doubt, or it could be external, such as peer disagreement. So it is fine for an evidentialist to say that Deanna (or Ankita if she is like her) should not necessarily conciliate, but at least re-open inquiry. If Ankita is like Efrosyni, then the evidentialist can't make that move. But she shouldn't want to. After all, Efrosyni has just as good reason to believe 22 times 18 is 396 as Deanna did after rechecking. So she has excellent reason to believe that 22 times 18 is 396. So she should keep believing it. It is much more plausible that Bojan made a rare mistake than that she made distinct mistakes on distinct calculations that ended up at the same point.

So once we realise that **Arithmetic** covers a range of possible cases, we see that the evidentialist can handle it. Indeed, they can arguably handle it better than the conciliationist can. But we don't need that to make the case. The best argument remaining for conciliationism was the argument from cases. And the best case for conciliationism was **Arithmetic**. If the evidentialist can handle even that case as well as the conciliationist, the evidentialist approach to disagreement is on very solid ground.

7 Appendix: Mathematical Evidence

It has been crucial to the story I've told that agents get evidence by thinking through problems. Efrosyni has more evidence than Deanna does for the proposition that 22 times 18 is 396. If that's right, the distinction between acquiring evidence and processing evidence is either much blurrier than epistemologists usually assume, or, more likely, is located in a very different place than epistemologists usually assume. Thinking through the problem is part of the evidence acquisition phase, not the evidence processing phase.

I'm going to argue for this conclusion by thinking through some examples involving mathematics. Let's start with a rather extreme example. A young mathe-

matics student, Tamati, is told about Fermat's Last Theorem. Tamati knows enough to understand exponentiation. So what he knows about basic arithmetic entails that Fermat's Last Theorem is true, though most people in his position would not believe this. (Indeed, it entails both that the Theorem is true, and that the usual statements of the theorem express truths.) Tamati is not told, and does not know, about the long history of Fermat's Last Theorem, and about how it puzzled mathematicians for centuries.

Tamati spontaneously forms the belief that Fermat's Last Theorem is true. This isn't because Tamati is usually prone to making mathematical guesses – this is the only time that he forms any such mathematical belief in such a spontaneous way. But nor is it because he is a mathematical savant; he can't articulate anything like a proof of Fermat's Last Theorem, or hint at what a proof of it would look like. To be sure, the ability to do this is no sign of ignorance. I know that Fermat's Last Theorem is true, but I don't know how the proof goes. Still, in Tamati's case, it does seem that he lacks knowledge.

And Tamati's lack of knowledge here is something that needs explaining. I think that the best explanation for it is that he lacks evidence. What other explanations are available?

Could it be that Tamati is unreliable? It doesn't seem so; by hypothesis he is perfectly reliable. Could it be that Tamati's method is unreliable? Again, it doesn't seem so. The method of spontaneously forming mathematical beliefs about mathematical propositions is one that we use every day, and usually reliable. It's true that if we restricted the scope of the methods to 'complex' propositions, it would be true that Tamati's method is unreliable. But I don't think there's any way of giving a sense in which Fermat's Last Theorem is complex without appeal to some independent epistemological notion. It isn't, after all, complex to state. What is complex is proving it. The key point here is that the only way to draw a reference class of exactly the right breadth to make Tamati's belief-forming method unreliable is to use other epistemological ideas, like the idea that Fermat's Last Theorem is hard to prove. That suggests it is the other epistemological ideas that explain Tamati's lack of knowledge, not the unreliability of his method relative to a somewhat gerrymandered reference class. And note that even if we get the reference class right, we have to somehow make it the case that the fact that Tamati is himself quite reliable when spontaneously forming mathematical beliefs, even about complex cases, irrelevant. I don't claim to have sealed the case, but I think it is going to be hard to explain Tamati's ignorance in terms of unreliability.

Could it be that, as Richard Fumerton (2010) says about a related case, that Tamati doesn't "appreciate" the evidence? I suspect that depending on how we understand appreciation, this explanation will either be too strong or too weak. If

appreciation means understanding how and why the evidence supports the conclusion, then too few beliefs end up being justified. Before they take a logic class, introductory students can come to know that Ga by inferring it from Fa and $\forall x(Fx \rightarrow Gx)$. But they don't need to know how or why their knowledge supports Ga . Indeed, they can be radically mistaken about the nature of that support, as many introductory logic students are, and still know Ga on that basis. On the other hand, if appreciation means having a true belief that the evidence supports the conclusion, we can posit that Tamati believes that. Tamati, we might imagine, knows enough to know that mathematical truths are entailed by any proposition, so he believes that his evidence entails Fermat's Last Theorem. That can't be enough for knowledge. If we insist that appreciation means knowing that the evidence supports the conclusion, then we have to explain why it is that Tamati doesn't know exactly this, and this doesn't seem any easier than the problem we started with.

So I think the best explanation of Tamati's ignorance is that he lacks sufficient evidence to know that Fermat's Last Theorem is true. When one works through a mathematical, or a logical, or a philosophical, problem, one gets evidence about the correct answer to that problem. But Tamati lacks that evidence. And he doesn't have any other form of evidence, like the testimonial evidence most of us have, that would make up for the lack of mathematical evidence. So thinking through a mathematical problem is not just a matter of processing evidence, it is a way of acquiring new mathematical evidence.

This way of thinking about mathematics is hardly radical. It is a commonplace in mathematics that one can get evidence for or against mathematical propositions. Philosophers too often think that evidence that entails a conclusion is maximally strong evidence. This assumption is encoded in most probabilistic models of evidence. But it isn't true. Facts about Andrew Wiles's diet are terrible evidence that Fermat's Last Theorem is true, even though they entail its truth. Facts about what he wrote in his notebooks, on the other hand, are excellent evidence that it is true. Thinking that entailing reasons are maximally strong reasons seems just to confuse inference with implication (Harman, 1986).

This attitude, of thinking that entailing reasons are maximally strong reasons, goes along with another bad attitude it is easy to adopt, namely when p is a mathematical proposition, that our evidence supports either a maximally strong belief in p or a maximally strong belief in $\neg p$. There seem to be numerous counterexamples to this view. Sanjoy Mahajan (2010) describes a lot of heuristics that can be used to quickly refute various mathematical hypotheses. The heuristics involve, for example, checking whether the 'dimensions' of a proposed identity are correct, checking limit cases, and that sort of thing. So consider the hypothesis that the area of an ellipse is πab , where a is the distance from the centre to the nearest point on

the ellipse, and b is the distance from the centre to the furthest point. After going through a number of other proposals and showing how they can be quickly refuted, Mahajan says this about the proposal that the area is πab .

This candidate passes all three tests. ...With every test that a candidate passes, confidence in it increases. So you can be confident in this candidate. And indeed it is correct. (Mahajan, 2010, 21)

It might be worried that the position I'm adopting here, that we often need evidence of a connection between premises and conclusion to infer the conclusion from the premises, even when the premises entail the conclusion, risks running into the regresses described by Lewis Carroll (1895). It certainly would be bad to say that to infer q from p and $p \rightarrow q$, the agent needs to know $(p \wedge (p \rightarrow q)) \rightarrow q$. That way lies regress, and perhaps madness. But that's not what is being claimed here. Rather, the claim is that for non-obvious entailments, the agent must know that the entailment obtains, or at least know the corresponding material implication, in order to use it in reasoning. It's consistent with that that we don't need any extra evidence to make simple inferential moves, either in deductive or inductive reasoning. And that's all we need to say to stop the regress.

Once we think about evidence this way, it becomes clear how rare it would be for people to have the same evidence. They would not just have to have been exposed to the same features of the external world, they would have had to perform the same evidence-gathering internal processes. If we set up puzzles about peer disagreement by stipulating that not only is that true, but that the parties know it to be true, we restrict the puzzles to a vanishingly small, and vanishingly unimportant, class. Happily, the evidentialist theory has something to say about a much broader class of cases than that.

References

- Ballantyne, Nathan and Coffman, E. J. 2011. "Uniqueness, Evidence and Rationality." *Philosophers' Imprint* 11:1–13.
- . 2012. "Conciliationism and Uniqueness." *Australasian Journal of Philosophy* 90:657–670, doi:10.1080/00048402.2011.627926.
- Carroll, Lewis. 1895. "What the Tortoise Said to Achilles." *Mind* 4:278–280.
- Christensen, David. 2007. "Epistemology of disagreement: The good news." *Philosophical Review* 116:187–217, doi:10.1215/00318108-2006-035.
- . 2009. "Disagreement as evidence: The epistemology of controversy." *Philosophy Compass* 4:756–767.

- . 2010. “Higher-order evidence.” *Philosophy and Phenomenological Research* 81:185–215, doi:10.1111/j.1933-1592.2010.00366.x.
- . 2011. “Disagreement, Question-Begging and Epistemic Self-Criticism.” *Philosophers’ Imprint* 11:1–22.
- . 2014. “Conciliation, Uniqueness and Rational Toxicity.” *Nous* , doi:10.1111/nous.12077 First published online 8 October 2014.
- Christensen, David and Lackey, Jennifer (eds.). 2013. *The Epistemology of Disagreement: New Essays*. Oxford: Oxford University Press.
- Cohen, Stewart. 2013. “A Defence of the (Almost) Equal Weight View.” In Christensen and Lackey (2013), 98–117.
- Douven, Igor. 2009. “Uniqueness Revisited.” *American Philosophical Quarterly* 46:347–361.
- Elga, Adam. 2007. “Reflection and disagreement.” *Nous* 41:478–502, doi:10.1111/j.1468-0068.2007.00656.x.
- Enoch, David. 2010. “Not Just a Truthometer: Taking Oneself Seriously (but not too seriously) in Cases of Peer Disagreement.” *Mind* Forthcoming.
- Feldman, Richard. 2007. “Reasonable religious disagreements.” In *Philosophers Without Gods: Meditations on Atheism and the Secular*, 194–214. Oxford: Oxford University Press.
- Fumerton, Richard. 2010. “You Can’t Trust a Philosopher.” In Warfield and Feldman (2010), 91–110.
- Harman, Gilbert. 1986. *Change in View*. Cambridge, MA: Bradford.
- Jackson, Frank. 1987. *Conditionals*. Blackwell: Oxford.
- Jehle, David and Fitelson, Branden. 2009. “What is the “Equal Weight View”?” *Episteme* 6:280–293, doi:10.3366/E1742360009000719.
- Kelly, Thomas. 2005. “The epistemic significance of disagreement.” *Oxford Studies in Epistemology* 1:167–196.
- . 2010. “Peer disagreement and higher order evidence.” In Warfield and Feldman (2010), 111–174.

- Lackey, Jennifer. 2010. "What Should We Do When We Disagree." *Oxford Studies in Epistemology* 3:274–293.
- Lam, Barry. 2011. "On the rationality of belief-invariance in light of Peer disagreement." *Philosophical Review* 120:207–245.
- Lasonen-Aarnio, Maria. 2013. "Disagreement and Evidential Attenuation." *Nous* 47:767–794.
- . 2014. "Higher-Order Evidence and the Limits of Defeat." *Philosophy and Phenomenological Research* 88:314–345.
- Levinstein, Ben. 2013. *Accuracy as Epistemic Utility*. Ph.D. thesis, Rutgers University.
- Loewer, Barry and Laddaga, Robert. 1985. "Destroying the Consensus." *Synthese* 62:79–95.
- Lord, Errol. 2014. "From Independence to Conciliationism: An Obituary." *Australasian Journal of Philosophy* 92:365–377, doi:10.1080/00048402.2013.829506.
- Mahajan, Sanjoy. 2010. *Street-Fighting Mathematics: The Art of Educated Guessing and Opportunistic Problem Solving*. Cambridge, MA: MIT Press, second edition.
- Moss, Sarah. 2011. "Scoring Rules and Epistemic Compromise." *Mind* 120:1053–1069, doi:10.1093/mind/fzs007.
- Polymath, D. H. J. 2014. "New equidistribution estimates of Zhang type, and bounded gaps between primes." Retrieved from <http://arxiv.org/abs/1402.0811>.
- Schoenfield, Miriam. 2014. "Permission to Believe: Why Permissivism Is True and What It Tells Us About Irrelevant Influences on Belief." *Nous* 48, doi:10.1111/nous.12006.
- Treynor, Jack L. 1987. "Market Efficiency and the Bean Jar Experiment." *Financial Analysts Journal* 43:50–53.
- Warfield, Ted and Feldman, Richard (eds.). 2010. *Disagreement*. Oxford: Oxford University Press.
- Weatherson, Brian. 2013. "Disagreements, Philosophical and Otherwise." In Christensen and Lackey (2013), 54–73.
- White, Roger. 2005. "Epistemic permissiveness." *Philosophical Perspectives* 19:445–459.

—. 2009. “On Treating Oneself and Others as Thermometers.” *Episteme* 6:233–250, doi:[10.3366/E1742360009000689](https://doi.org/10.3366/E1742360009000689).

Wittgenstein, Ludwig. 1953. *Philosophical Investigations*. London: Macmillan.

Worsnip, Alex. 2014. “Disagreement about disagreement? What disagreement about disagreement?” *Philosophers’ Imprint* 14:1–20.