

Disagreement

ABSTRACT

There has been a recent explosion of interest in the epistemology of disagreement. Much of the recent literature is concerned with a particular range of puzzle cases (discussed in the first section of my paper). Almost all of the papers that contribute to that recent literature make mention of questions about religious disagreement in ways that suggest that there are interesting connections between those puzzle cases and real life cases of religious disagreement. One important aim of my paper is to cast doubt on that suggestion. More generally, the aim of my paper is to give a reasonably full account of the recent literature on the epistemology of disagreement, and then to give a serious discussion of some of the epistemological issues that are raised by real world religious disagreements.

KEYWORDS: disagreement; doxastic peer; epistemology of disagreement; reasonable disagreement; religious disagreement; shared evidence

1. Cases

Here are some cases of peer disagreement in which nearly everyone supposes that it is simply wrong to hold fast—i.e., to fail to adjust credences to bring them closer to the credences of another.

Restaurant: A number of philosophers go out to dinner after a seminar. When it comes time to pay the bill, they all agree to a 15% tip (plus rounding), and to split the whole cost evenly between them. With the bill clearly visible to all, two of the members of the party calculate the individual shares in the bill. One calculates \$63, and, at this point, is confident that each owes \$63. The other then tells the first she has calculated \$61. How should the first calculator react to this new information? We can add to the story that it is mutual knowledge that the group has a long history of eating out together; that the two in question always calculate the bill; that they're equally reliable at performing this calculation—mostly they agree, but when they don't, neither is more likely than the other to have things correct; that neither is especially tired or energetic; neither has consumed more coffee or wine than the other; neither has seemed particularly dull or sharp over the course of the evening; neither feels more confident than the other about the result of this particular calculation; and so forth. Given enough of this type of filling out of the story, it seems inescapable that the first should reduce his confidence that each owes \$63 when he learns of the calculation of the other.

Wrist Watch: Identical twin children have been given new watches for their birthday. They synchronise their watches before they go to bed. The next morning, one of the twins awakes and sees that her watch reads 6:45 a.m.; at this point, she is confident that the time is 6:45 a.m. Then the other twin awakes, and announces that her new watch says that the time is 6:51 a.m. How should the first twin react to this new information? We can add to the story further details that make it mutually known that it is incredibly unlikely that the second twin is lying, or misreading her watch, and the like. Given enough of this type of filling out of the story, it seems inescapable that the first twin should reduce her confidence that it was 6:45 a.m. when she first looked at

her watch (and should give increased confidence to the proposition that one of the new watches is not working properly).

Horse Race: Avid punters are gathered at the finishing line of the last at Caulfield. Two horses flash past the finishing line, with very little separating them. One of the punters leaps into the air, shouting that the horse that he has backed—No. 4—was first past the post. At this point, the punter in question is very confident that he has backed the winning horse. However, he immediately becomes aware that another of the punters has leapt into the air, shouting that the horse that he has backed—No.8—was first past the post. How should the first punter react when he obtains this new information? As in the previous cases, we can fill in details to the story—concerning the mutual knowledge of the two punters that each is no more likely than the other to have correctly determined the winning horse—that make it seem inescapable that the first punter should reduce his confidence that he has backed the winning horse when he observes the behaviour of the second punter.

Perfect Pitch: Two musicians hear a note played on a violin. One judges that the note is a C. At this point, the musician is very confident that the note is a C. She then learns that the other musician thinks that the note was a C-sharp. How should the first musician react when she obtains this new information? As in the previous cases, we can fill in the details of the story—concerning the mutual knowledge of the two musicians that each is no more likely than the other to have correctly determined the pitch of the played note—that make it seem inescapable that the first musician should reduce her confidence that the note that she heard was a C when she learns of the judgment of the second musician.

There are a number of features that these cases have in common. In particular, these are all cases in which it is plausible to suppose that, no matter what views they might have had some time ago about their relative abilities on the domains in question, the people involved could have collected enough uncontroversial data to bring them to mutual knowledge that they are peers with respect to success at the relatively encapsulated task that is at issue: performing simple arithmetic operations, detecting the winners of horse races by visual observation, determining the time by reading a watch, identifying the pitch of a note simply by listening to it, and so forth. Moreover, these are also all cases in which it is plausible to suppose that, no matter what views they might have had some time ago about their relative abilities on the domains in question, the people involved could have collected enough uncontroversial data to bring them to mutual knowledge that it is quite common even for experts to make errors on the tasks in question, and even in quite friendly conditions: even mathematical experts do sometimes make errors in adding up columns of figures; even expert race callers do sometimes call the wrong horse; even brand new watches do sometimes fail to function properly; even those with perfect pitch sometimes get the identification of the pitch of a note wrong.

Here are some cases of peer disagreement in which nearly everyone supposes that it is perfectly acceptable for one to hold fast—i.e. to fail to adjust credences to bring them closer to the credences of another.

Elementary Math: Two people who have been colleagues for the past decade are drinking coffee at a café while trying to determine how many other people from their

department will be attending an upcoming conference. One, reasoning aloud, says: ‘Well, John and Monima are going on Wednesday, and Karen and Jakob are going on Thursday, and, since $2+2=4$, there will be four other people from our department at the conference’. In response, the other says: ‘But $2+2$ does not equal 4’. Prior to the disagreement, neither party to the conversation has any reason to suppose that the other is evidentially or cognitively deficient in any way; and, we may suppose, each knows that none of the speech acts involved is insincere. Moreover, we may suppose, the one is feeling perfectly fine: the one has no reason to think that she is depressed, or delusional, or drugged, or drunk, and so forth. In this case, it seems plausible to suppose that the one should conclude that something has just now gone evidentially or cognitively awry with the second, and that the one should not even slightly adjust the credence she gives to the claim that $2+2=4$.

Perception: Three people who have been roommates for the past decade are eating lunch together in the kitchen in their apartment. One asks a second to pass the parmesan cheese to the third, whereupon the second announces that the third is not present. Prior to the announcement, neither the first nor the second has had any reason to think that the other is evidentially or cognitively deficient in any way; and, we may suppose, each knows that none of the speech acts involved is insincere. Moreover, we may suppose, the one is feeling perfectly fine: the one has no reason to think that she is depressed, or delusional, or drugged, or drunk, and so forth. In this case, it seems plausible to suppose that the one should conclude that something has just now gone evidentially or cognitively awry with the second, and that the one should not even slightly adjust the credence she gives to the claim that the third is sitting and eating in the dining room.

Directions: I have lived in Glen Waverley all my life, and I am very familiar with the layout of streets in my neighbourhood. I regularly dine at the Rose of Thailand on Coleman Parade. My next door neighbour, Pat, has also lived in Glen Waverley all his life, and he also dines often at the Rose of Thailand; indeed, it is not uncommon for me to bump into him at that establishment. Today, when I bump into him as I pass his front door, I tell him that I’m on my way to the Rose of Thailand. He says: ‘But the Pinewood shops are in the other direction.’ I say: ‘The Rose of Thailand is on Coleman Parade’; and he replies: ‘No, the Rose of Thailand is not on Coleman Parade—it’s at the Pinewood shops.’ Prior to this little exchange, neither I nor Pat has had any reason to think that the other is evidentially or cognitively deficient in any way; and, we may suppose, we both know that none of the speech acts involved is insincere. Moreover, I am feeling perfectly fine: I have no reason to think that I am depressed, or delusional, or drugged, or drunk, and so forth. In this case, too, it seems plausible to suppose that I should conclude that something has just now gone evidentially or cognitively awry with Pat, and that I should not even slightly adjust the credence I give to the claim that the Rose of Thailand is on Coleman Parade.

Bird: While reading in the library with your best friend, you glance out the window, catch a glimpse of a bird flying by, and on this basis form the belief that a magpie just flew by. After saying to your friend who was looking out the window at the same time, that you enjoyed seeing the magpie that just flew by, your friend responds: ‘Nothing flew by the window’. Prior to this episode, neither you nor your friend had any reason to think that the other is evidentially or cognitively deficient in any way, and, we may suppose, each knows that none of the speech acts involved is insincere. Moreover, we

may suppose, you are feeling perfectly fine: you have no reason to think that you are depressed, or delusional, or drugged, or drunk, and so forth. In this case, yet again, it seems plausible to suppose that you should conclude that something has just now gone evidentially or cognitively awry with your friend, and that you should not even slightly adjust the credence that you give to the claim that a magpie flew past the window.

There are a number of features that these cases have in common. In particular, in each of these cases, an erstwhile peer evinces a judgment that conflicts with another's cognitively basic judgments: judgments *immediately grounded* in memory, or perception, or elementary arithmetic, or the like. Of course, it is also true that, in each of these cases, an erstwhile peer evinces a judgment that conflicts with another's very confident and highly justified beliefs—but it is worth noting from the beginning that what seems to drive intuitions about these cases is the basic nature of the cognitive judgements that they involve. How could any erstwhile epistemic peer be confused about whether $2+2=4$ unless in the grip of some phenomenologically detectable cognitive misfortune? How could any erstwhile epistemic peer be unable to see what is right in front of his eyes unless in the grip of some phenomenologically detectable cognitive misfortune? How could any erstwhile epistemic peer be unable to recall uncontroversial and well-entrenched material from memory unless in the grip of some phenomenologically detectable cognitive misfortune? Etc.

2. Key Concepts

In order to draw lessons from the above cases, there are various concepts that we need to invoke. Some of these concepts are employed in the setting up of the above cases; other concepts are really only needed in the analysis of the cases.

Disagreement: People *disagree* when they have different doxastic attitudes towards a proposition. Thus, for example: (i) two people disagree about the proposition that p if one believes that p and the other believes that not p ; (ii) two people disagree about the proposition that p if one believes that p and the other fails to believe that p ; and (iii) two people disagree about the proposition that p if they give different credences to the claim that p (e.g. if one holds that it is very probable that p , and the other holds that it is only just more probable than not that p). People have a *disagreement* about a proposition when they disagree about that proposition, and it is mutual knowledge between them that they disagree about that proposition. Of course, when people have a disagreement about a proposition, those people typically have a disagreement about many propositions. In particular, if people are rational, and they have a disagreement about a proposition, then those people will typically have a disagreement about other propositions. For example, if A and B have a disagreement about whether that p , and if A and B have an agreement about whether that q , then, at least in the typical case, A and B will also have a disagreement about whether that p and q .

Reasonable Disagreement: People *reasonably disagree* about a proposition when they disagree about that proposition, and they are reasonable in their doxastic attitude towards that proposition. People have a *reasonable disagreement* about a proposition when they reasonably disagree about that proposition, and it is mutual knowledge between them that they disagree about that proposition. People *reasonably agree to disagree* about a proposition when they have a reasonable disagreement about that

proposition in circumstances in which each believes that those with whom they disagree are also reasonable in their doxastic attitude toward the proposition in question. Of course, we could go on to make further definitions here: we could, for example, have a term to cover cases of reasonable disagreement about a proposition in circumstances in which each *reasonably* believes that those with whom they disagree are also reasonable in their doxastic attitude toward that proposition. But we should not multiply definitions beyond those that we will actually use in our discussion.

Cognitive Comparators: People are *cognitive peers* if they are cognitive equals: equally intelligent, equally adept at reasoning, equally good at memory recall, and so forth. One person is *cognitively superior* to another if the one is superior to the other across all cognitive domains: more intelligent, more adept at reasoning, better at memory recall, and so forth. (Perhaps we might also want to say that one person is cognitively superior to another if the one comes out ahead of the other on some appropriate weighting of comparisons on particular cognitive domains. However, there is no need for us to pursue that line of thought here.) If there are domain-specific cognitive skills, then we can relativise cognitive comparisons to domains of inquiry: we can say, for example, that people are *cognitive peers with respect to a given domain of inquiry* just in case they are cognitive equals on that domain of inquiry. Subsequent discussion is simplified if we ignore the possibility that there are domain-specific cognitive skills.

Evidential Comparators: People are *evidential peers* if they are evidential equals, i.e. if they are equally well-informed. Since informedness is clearly domain-specific, we do need here to relativise to domains of inquiry (and perhaps even to particular propositions): people are *evidential peers with respect to a given domain of inquiry* just in case they are equally well-informed about that domain of inquiry. If we suppose that it is possible for people to fully share their evidence, then we might suppose that one way in which people could come to be evidential peers is through full sharing of evidence. However, on plausible conceptions of evidence, it is not possible for people to fully share their evidence—i.e. to fully exchange evidence through discussion. On the one hand, on plausible conceptions of evidence, people do not have the capacity to articulate all of the evidence that they currently possess. And, on the other hand, on plausible conceptions of evidence, people may well no longer be in possession of some of the total evidence upon which their views are founded. There are good grounds for thinking that it is no requirement of reason that people be able to keep track of the grounds for their beliefs: owing to limited storage capacity and so forth, people often quite properly form beliefs and then discard the evidence that constitutes the basis for their formation. (See Harman (1986) for further discussion of this kind of point.)

Doxastic Comparators: People are *doxastic peers* if they are both cognitive peers and evidential peers. Doxastic peers are equals in respect of both informedness and cognitive skills: intelligence, reasoning, memory recall, and so forth. It is standard in the literature to refer to doxastic peers as ‘epistemic peers’. However, it is clearly a controversial question—affected, for instance, by disputes between internalists and externalists about the nature of doxastic justification—whether, other things being equal, doxastic peers who believe alike also know alike. Some contributors to the literature—e.g. Elga (2007a)—have insisted on a conception of doxastic peers that requires equal probability of formation or possession of true belief: A counts B as a

doxastic peer with respect to an about-to-be-judged claim *iff* A thinks that, conditional on A and B having a disagreement about that claim, they are equally likely to be mistaken. However, Elga's formulation greatly reduces the range of applicability of the notion of a doxastic peer, since it is simply silent on cases in which A and B assign prior probabilities to claims of interest, i.e. on cases in which A and B have already made judgments about a claim. And, of course, Elga's formulation won't admit of interesting extension to cases of already-judged claims, because A can only judge that B is equally likely to be mistaken if A supposes that B makes the *same* probability judgment that A makes.

Belief: There are two different ways in which we might think about belief. On the one hand, we might think about belief as an *all-or-nothing* attitude: for any proposition that p, either one believes that p, or one believes that not-p, or one suspends judgment about whether that p, or one simply has no attitude at all towards whether that p. On the other hand, we might think about belief as a *graded* attitude, properly represented by a credence function, or a family of credence functions, or the like: for any proposition that p, one assigns a particular credence—or a particular range of credences, or the like—to that proposition, or else one has no attitude at all towards that proposition. While some of the literature on the epistemology of disagreement is couched in terms of all-or-nothing doxastic attitude—see, for example, Feldman (2007)—it seems to me desirable to couch discussion in the more general framework of graded doxastic attitudes.

Reason: It is relatively easy to say what the sense of 'reason'—'reasonable', 'rational'—that is operative in discussion of the epistemology of disagreement is *not*. As Feldman (2007) observes, we are not meaning to talk about general tendencies to thought and behaviour, or about escapes from flagrant unreason, or about prudential values inherent in beliefs, or about features of behavioural choices. What we do have in mind is some *normative* standard to which belief revision and belief formation ought to conform. However, it is no easy matter to delineate this standard. Feldman thinks that what is at issue is 'adequate evidential support'. But it seems to me that considerations about evidence are only one part of the canon that governs belief revision and belief formation. For now, I shall simply suppose that we do have a satisfactory shared conception of the norms that govern belief revision and belief formation; I think that it is a major theoretical undertaking to give an adequate account of the relevant norms. I shall further suppose that we have a good understanding of some of the ways in which these norms can be violated: in particular, we are reasonably good at recognising some kinds of cases of *dogmatism*—where subjects fail to revise their beliefs in line with the beliefs of others in circumstances in which they ought thus to revise—and some kinds of cases of *servility*—where subjects revise their beliefs in line with the beliefs of others in circumstances in which they ought not thus to revise.

3. Positions

Much of the very recent literature on the epistemology of disagreement is couched in terms of two extreme general positions. We are to imagine that A and B are doxastic peers [on a given domain], that they take each other to be doxastic peers [on that domain], that they form independent judgments about a particular proposition that p [belonging to the domain in question], and that they then *learn* that they disagree in

their judgment about that proposition. The question we are to answer is: how should A and B adjust their credences when they learn that they disagree about whether that p?

1. The *conformist* or *conciliationist* answer to our question is that, in every case, A and B should each adjust his or her credence to bring it closer to the credence of the other party. Perhaps A and B should *split the difference* between their views; perhaps A and B should give *equal weight* to the view that they previously held and the view that the other person holds; perhaps A and B should somehow give *proportional weight* to the view that they previously held and the view that the other person holds (perhaps depending in some way upon the *total evidence* that each has).

2. The *non-conformist* or *steadfast* answer to our question is that, in every case, there is no requirement that either A or B adjust his or her credence to bring it closer to the credence of the other party. It may be that each of A and B should give *no weight* to the view of the other party, perhaps on the grounds that he or she possesses some *insight* that the other fails to have. Or it may be that, if one of A or B has *right reasons* for holding his or her own view while the other party does not, then the one ought not to revise his or her view in the face of his or her peer's dissenting judgment.

Given the cases introduced in section 1 above, it seems that neither of these extreme general positions can be correct. In the cases of *Elementary Math*, *Perception*, *Directions* and *Bird*, almost everyone supposes that it is perfectly acceptable for the one to hold fast—i.e. to fail to adjust credences to bring them closer to the credences of the other. However, in the cases of *Restaurant*, *Wrist Watch*, *Horse Race* and *Perfect Pitch*, almost everyone supposes that it is simply wrong for the one to hold fast—i.e., to fail to adjust credences to bring them closer to the credences of the other.

An obvious response to the failure of the two extreme positions is to look for an intermediate view. Those who are broadly sympathetic to the conformist or conciliationist view might suggest that what is needed is a clause that removes the obligation to shift credence in cases in which the one has good reason to—at least temporarily—retreat from the view that the other is a doxastic peer (as, plausibly, in the cases of *Elementary Math*, *Perception*, *Directions* and *Bird*). Those who are broadly sympathetic to the non-conformist or non-conciliationist view might suggest that what is needed is a clause that introduces an obligation to shift credences in cases in which one's total evidence gives no stronger support to one's own view than it gives to the view of one's peer (as, plausibly, in the cases of *Restaurant*, *Wrist Watch*, *Horse Race* and *Perfect Pitch*). Those who have no particular sympathy for either of the extreme positions might look for a more evidently mixed view: perhaps, for example, a view that says: (a) if one has formed a very confident belief that is highly justified, then one need not adjust credence exactly if one feels oneself to be in good doxastic shape; and (b) if one has formed a belief that is either insufficiently confident or insufficiently well-justified, then one is required to adjust one's credence in the direction of one's peer (cf. Lackey (2009a)(2009b)).

There are at least two different ways in which these kinds of responses might be criticised. On the one hand, it might be thought that there are more general theoretical standpoints that rule out all such responses. Thus, for example, Bayesians will say that there is a simple, general answer to the question of how to respond to evidence: one

ought always to conditionalise! Simple inspection quickly reveals, I think, that none of the proposals that we have canvassed to date is consistent with this Bayesian advice. (See Gardner and Oppy (in preparation).) On the other hand, it might be thought that a better response to the failure of the two extreme positions is to reconsider the question that those positions are intended to address. In particular, one might choose to reconsider the implicit assumption that there is something interesting that can be said at this level of generality about this particular kind of doxastic disagreement. (Of course, it is also open to contest the responses on their own terms, i.e. to look for counterexamples, and so forth. Thus, for instance, one might think that the last of the proposals mentioned above cannot handle the following case. *Extreme Horse Race*: We've judged 20,000 horse races, and agreed every time, and been right every time, and we both know this. This time, we disagree. I (justifiably) feel myself to be in fine doxastic shape. Prior to learning that we disagree, I have a very high degree of justified confidence in my judgment. But surely the fact that I'm (justifiably) feeling in good doxastic shape is not enough to justify my simply ignoring your conflicting judgment.)

Suppose that we do arrive at an answer to our initial question. That is, suppose that we do have a prescription governing the revision of the credences of A and B in circumstances in which A and B are doxastic peers [on a given domain] who take each other to be doxastic peers [on that domain], who independently form judgments about a particular proposition that p [belonging to the domain in question], and who then *learn* that they disagree in their judgment about that proposition. Is there any reason to suppose that the answer to our initial question will generalise to cover other—more interesting—cases of peer disagreement (such as disagreements about religion, or politics, or the like)?

I don't think so. First, in many of the most interesting cases, the knowledge that one's opinions are contested by peers *antedates* the formation of those opinions. (This is most obviously true in philosophy: I continue to arrive at new philosophical views even though I know that any philosophical views that I develop are controversial.) Learning that *particular* peers disagree with me seems to be typically and properly of little moment. Second, in many of the most interesting cases, there is nothing that one can point to that constitutes *the* evidence and reasons that led one to acquire the particular opinions that one holds (nor, in many cases, is there anything that one can point to that constitutes *the* evidence and reasons that justify one in continuing to maintain the particular opinions that one holds). Third, as we noted earlier, in most interesting cases, it makes no sense to talk about total sharing of evidence; the most that we can reasonably hope to claim is that people are *roughly* equally well-informed about particular questions. Fourth, in many interesting cases, it is very hard to know what to make of the requirement that the judgments that we have formed are independent. In cases like *Restaurant*, *Wrist Watch*, *Horse Race* and *Perfect Pitch*, it is easy to see that the key judgments are independent; but how are we to assess the claim that the beliefs of two who share particular philosophical—or religious, or political—views are independent?

Further difficulties that arise here can be brought out by considering toy cases that lack some of the features of the test cases introduced earlier. Consider, for example:

Weather Forecast: Two meteorologists at the Weather Bureau are making predictions about the likelihood of rain tomorrow. The meteorologists make a thorough study of the current weather data and apply the various models that they use in order to make their predictions. One of the meteorologists forms the view that the probability of rain tomorrow is 55%; the other meteorologist forms the view that the probability of rain tomorrow is 35%. If we suppose that each meteorologist is aware of the judgment of the other—and if we also suppose that each regards the other as a doxastic peer—should we also suppose that each ought to move towards the level of confidence of the other?

Detective: Two very experienced and highly successful detectives—Dalglish and Morse—have been investigating a crime. On the basis of their joint investigation, Dalglish thinks that Mr. Big is the culprit, whereas Morse thinks that the crime is properly attributed to Dr. No. In the matter under investigation, some of the accumulated evidence points to each of the accused; and it is a very delicate matter to arrive at an all-things-considered judgment. If we suppose that each detective is aware of the judgment of the other—and if we also suppose that each regards the other as a doxastic peer—should we also suppose that each ought to move towards the level of confidence of the other?

Diagnosis: Dr. Haggard and Dr. Wolf are both impeccably educated and highly respected physicians with fifteen years of practising medicine under their belts. Ruby, who has been thoroughly examined by both doctors, has been suffering for months from extreme fatigue, swollen joints, muscle pain, memory loss, sensitivity to light, and persistent fevers. After running the same tests and receiving identical results, Dr. Haggard is quite confident that Ruby has chronic fatigue syndrome, and Dr. Wolf is just as confident that she suffers from lupus. If we suppose that each physician is aware of the judgment of the other—and if we suppose that each regards the other as a doxastic peer—should we also suppose that each ought to move towards the confidence level of the other?

In these cases, we are not considering relatively encapsulated tasks, and nor are we concerned with relatively basic cognitive judgments. On the contrary, we are considering complicated expert judgments that are grounded in protracted experience in the application of a diverse array of cognitive skills and capacities to the kind of question that is at issue. Because the judgments are grounded in protracted experience, there is a sense in which the two parties draw on very different bodies of evidence. True, the two parties may be able to share “the evidence” that belongs to the case in question: test results, lab reports, and so forth. But there is no way that they can share the rest of “the evidence” that contributes to the judgments that they make. (Apart from anything else, much of that evidence will simply have been forgotten.) However, that the two parties cannot share much of “the evidence” that contributes to their judgments does not yield a good reason for supposing that the two parties are not properly classified as doxastic peers: on the contrary, if we suppose that they are cognitive peers, that they have the same amount of experience, and that they have shared “the evidence” that belongs to the case in question—the test results, lab reports, and so forth—then we have a perfectly good sense in which the two parties are doxastic peers. To suppose otherwise would be to restrict questions about peer disagreement to a very narrow and relatively uninteresting class of cases.

In cases like *Weather Forecast*, *Detective* and *Diagnosis*, there is, I think, no simple argument for the conclusion that the protagonists are rationally *required* to adjust their credences when they learn about the conflicting judgments of other equally qualified parties. However, there is a relatively simple argument that suggests that it might be rationally *permissible* for the protagonists to fail to adjust their credences when they learn about the conflicting judgments of other equally qualified parties.

Suppose that, on a particular kind of judgment, we have two judges who follow methods that lead them to judge correctly in 95% of cases, but who get a different 5% of cases wrong. That is, in 90% of cases, they agree on the correct judgment; but, in 10% of cases, they disagree, and each is correct exactly half of the time. Suppose that we wish to use the judgments of these two judges—and nothing else but the judgements of these two judges—to arrive at a judgment of our own. How will we best proceed? Well, clearly enough, we should follow the lead of the two judges in the cases in which they agree: that gives us the correct answer 90% of the time. But, if we suspend judgment on the cases in which the judges disagree, we'll end up with a lower success rate than either of the judges does separately. Of course, where the judges are right 95% of the time and wrong 5% of the time, we'll be right 90% of the time, and suspending judgment 10% of the time. Will we be doing better? Well, it seems to me that that depends upon what you care about. There are at least two dimensions to good doxastic performance: avoiding error and gaining truth. Unless you think that *any* increase in avoidance of falsehood trumps *any* decrease in acquisition of truth, you're bound to think that there can be cases in which you do better to follow a single judge of the kind under consideration than to follow the combined strategy that requires suspension of judgment when the judges disagree. (Suppose, for example, that there are 20 such judges, each getting a different 5% of the cases wrong, and that you use the same strategy taking all 20 judges into account. Then you'll be suspending judgment 100% of the time. Are you doing better than a judge who gets 95% of cases right, and 5% of cases wrong?)

Suppose, now, that I am one of twenty judges, each of whom is a doxastic peer, and each of whom gets a different 5% of cases wrong. If I stick to my guns, I'll be getting 95% of cases right, and 5% of cases wrong. If I follow the strategy of suspending judgment whenever one of my doxastic peers disagrees, then I'll suspend judgment in every case: I'll get 0% of cases right and 0% of cases wrong. It seems to me that there is at least some reason to suppose that I will be doing worse if I adopt the latter strategy. But, at the very least, that suggests that there is a *prima facie* case for the permissibility of failure to adjust credences in at least some cases of peer disagreement involving expert judgment.

4. "Real" Cases

The recent literature on the epistemology of disagreement provides a wide range of suggestions for dealing with the "real" or "hard" cases—say, disagreements about religion, or politics, or the like. Some of these suggestions are broached against the background of commitment to a conformist—or non-conformist—line on the cases of peer disagreement canvassed in the previous parts of this paper. Here, however, I propose simply to consider each of these proposals on its own terms.

Hard Line: Feldman (2007) canvasses—but does not endorse—the view that there are no cases of reasonable peer disagreement. On his version of this line, all disagreement traces back either to unequal possession of information, or else to irrationality on the part of at least all sides but one in any dispute. Thus, when I find myself in persistent disagreement with the views of others, I am required to think either that they lack information that I possess, or else that they are irrational in their beliefs. Feldman claims that the hard line is very difficult to accept, not least in cases in which people have been carefully reflective and have openly discussed their differing responses. However, it is not so clear to me that this is so. As I noted earlier, a sensible definition of ‘doxastic peer’ will allow that doxastic peers differ in the evidence that they possess. Consequently, it seems plausible to suppose that it is actually true that, in ‘real’ cases of disagreement, there are relevant differences in the evidence possessed by parties to the disagreement. But, if this is right, then perhaps there can be reasonable peer disagreement in ‘real’ cases simply in virtue of differences in the evidence possessed or relied upon by parties to the disagreement.

Insight: Many authors—including van Inwagen (1995) and Plantinga (1995)—have supposed that, in cases of peer disagreement, I can only maintain that my own position is reasonable if I suppose that I have some kind of insight that other parties to the disagreement lack. On this line, I think, there can be cases of reasonable peer disagreement—but only because people can reasonably yet mistakenly believe that they possess insight that others do not have. For reasons given earlier, I do not think that appeal to insight is necessary to justify sticking to one’s guns in cases of peer disagreement. Others—e.g. Feldman (2007)—have argued that appeal to insight is not sufficient to justify sticking to one’s guns in cases of peer disagreement. Granted that each party to the dispute takes himself to have insight, what reason could any one of the parties have to privilege his own purported insight above the purported insights of his peers? I am not sure what to say about this; it may be that this debate will also simply degenerate into a dispute about the proper characterisation of doxastic peers.

Extreme Philosophical Scepticism: One extreme reaction to the possibility of peer disagreement would be to insist that all reasonable parties will withhold assent from any proposition about which there is the slightest possibility of peer disagreement. Perhaps one might suspect that this is one of the motivations that lay behind Pyrrhonian scepticism: if one never holds an opinion about any matter that is potentially controversial, then one will certainly avoid any of the possible pain and suffering that can come from being in states of conflict and disagreement. However, there are at least two obvious points to make in response to this suggestion. On the one hand, it is clearly extreme to give up any chance of arriving at true beliefs in order to avoid merely possible conflict with doxastic peers: surely one would do better to following a policy of withdrawing assent just in cases in which there is actual peer disagreement. And, in any case, on the other hand, it is not clear that extreme philosophical scepticism is a sustainable view: there are various reasons for thinking that it is self-undermining, or psychologically impossible, or the like.

Moderate Philosophical Scepticism: Some—e.g. Feldman (2007)—have defended the claim that all reasonable parties should withhold assent from any proposition about which there is actually peer disagreement. That is: in any actual case in which doxastic peers disagree about whether that *p*, those doxastic peers ought all to suspend judgment about whether that *p*. Feldman thinks that, while moderate philosophical

scepticism implies that many of us should give up many of our deeply held, well-considered beliefs, this conclusion is not distressing: “It calls for a kind of humility in response to the hard questions about which people so often find themselves in disagreement. It requires us to admit that we really do not know what the truth is in these cases.” (213) However, if—as I suggested above—we *can* reasonably agree to disagree, then there are ways of being humble that do not require that we give up our deeply held and well-considered beliefs. (In order to ward off worries about self-defeat, perhaps we can follow Elga (2007b) in claiming that fundamental doxastic principles are not required to apply to themselves: the moderate philosophical sceptic is permitted not to withhold assent from the claim that *in any actual case in which doxastic peers disagree about whether that p, those doxastic peers ought all to suspend judgment about whether that p* even in the face of peer disagreement about that claim.)

Merely Verbal Disagreement: Sosa (unpublished) argues that “Perhaps the main reason why disagreements on controversial issues are often reasonable on all sides is that they are to some extent verbal. ... Disagreement can be reasonable and interesting when its nature as verbal can emerge only through patient, extended analysis and discussion.” (4) I agree that it can be unclear whether philosophical disputes are largely verbal. Some people think that the dispute between three-dimensionalists and four-dimensionalists is primarily a matter of book-keeping. Other people think much the same about the dispute between neo-Millians and neo-Fregeans on the semantics of propositional attitude ascriptions. However, I can’t help feeling that, if these people are right, and those disputes do turn out to be largely verbal, then it also turns out that those disputes are not *really* interesting. It is not even *prima facie* plausible to suppose that very large debates in religion, politics, and so forth are significantly merely verbal (and the same can also be said for cases like *Weather Forecast*, *Diagnosis* and *Detective*). Even if there are some cases in which disagreements on controversial issues are reasonable on all sides because they are to some extent verbal, those cases are not the ones that deserve serious attention.

Inscrutability of Reasons: A reason that some might give for supposing that doxastic peers can reasonably agree to disagree is that reasons and evidence can be inscrutable. The reasons why we believe as we do are often not transparent to us, and cannot be rendered fully transparent to us by any amount of further inquiry; the evidence that supports the formation of our beliefs is progressively discarded or otherwise forgotten. Moreover, the reasons that *other* people have for believing as they are do are less transparent to us than our own reasons, and even more resistant to fully transparent rendition by subsequent inquiry. Given these considerations about the inscrutability of reasons and evidence, surely we have good reason to be cautious in adjusting our beliefs in the face of mere peer disagreement. Or do we? If it is true that reasons and evidence can be inscrutable, then the consequences will be the same for everyone: that fact gives me no reason at all to suppose that my beliefs are more likely to be true than yours—so how can it be that that fact could support the idea that doxastic peers can reasonably agree to disagree? Well, perhaps like this: While it is true that I have no good reason for thinking that I have more true beliefs than you, I do have good reason for thinking that, were I to suspend judgment on all the points on which we disagree, I would then have fewer true beliefs than you, unless you also followed me in suspending judgment on the points on which we disagree. I have no reason to suppose that I would do better by taking on your beliefs—and there would clearly be

costs involved. However, I have good reason to suppose that I would do worse were I to suspend judgment where we disagree. So, I should stay where I am: we should agree to disagree!

Difference in Starting Points: A related reason that some might give for supposing that doxastic peers can reasonably agree to disagree is that we have differing prior probabilities (or differing presuppositions, or differing worldviews, or differing starting assumptions, or differing doxastic frameworks, or the like). Given these considerations about difference of starting points, surely we have good reason to be cautious in adjusting our beliefs in the face of mere peer disagreement. Or do we? As before, if it is true that we have different doxastic starting points, then the consequences are the same for everyone: that fact gives me no reason at all to suppose that my beliefs are more likely to be true than yours—so how can it be that that fact supports the idea that doxastic peers can reasonably agree to disagree? (See Feldman (2007: 205f.) for a similar line of thought.) But, given the line that I ran in the previous case, it should be clear what I think might be said in response. If we think of our doxastic peers as those who are getting about as much of the truth as we are, and if what we care about is having the best ratios that we can of truth to falsehood and truth to suspension of judgment, and if we think that we and our doxastic peers are doing reasonably well in obtaining truth, then we actually have good reason to stick to our guns: we will not *improve* our views by suspending judgment where we disagree.

Insufficient Agreement: Another related reason that some might give for supposing that doxastic peers can reasonably agree to disagree is that, in interesting cases of ‘real’ peer disagreement, there is disagreement about a wide range of more or less related claims. Even in cases like *Weather Forecast*, *Detective*, and *Diagnosis*, it is not the case that the protagonists simply disagree about the judgment that is at issue. Rather, their disagreement on the judgment at issue is a manifestation of their disagreement on a range of related questions. And, in ‘real’ cases, this further disagreement is typically even more extensive. But, given these considerations about widespread differences in belief, surely we have good reason to be cautious in adjusting our beliefs in the face of mere peer disagreement. Or do we? As before, if it is true that there is widespread difference in belief, then the consequences are the same for everyone: that fact gives me no reason at all to suppose that my beliefs are more likely to be true than yours—so how can it be that that fact supports the idea that doxastic peers can reasonably agree to disagree? Perhaps there is a practical problem in determining exactly how much one should stop believing in the face of peer disagreement on a particular issue—but, in principle, it might seem that one ought to withhold judgment wherever one has a doxastic peer who makes a dissenting judgment. (Elga 2007a) suggests that in cases in which there is insufficient agreement, there are good grounds for rejecting claims to doxastic peerhood. But, as I noted earlier, he works with a very restrictive notion of doxastic peerhood.) Of course, I think that this case is vulnerable to the same kind of response as I gave in the previous two cases.

Alethic Impurity: A final, rather different, reason for supposing that doxastic peers can reasonably agree to disagree is that, in interesting cases of ‘real’ peer disagreement, it is very hard to disentangle clashes in alethic attitudes from clashes in non-alethic attitudes. For instance, in disagreements about ethics and politics, it is not a straightforward matter to determine to what extent those disagreements concern

desires rather than beliefs. But, given that there are uncertainties about the extent to which the disagreements concern desires rather than beliefs, there are reasons to resist taking a conciliatory approach to these disagreements. (That you want something other than what I want does not give me a reason to give up my own desire and take on yours.) Note that this suggestion does not involve a commitment to non-cognitivism about the matters in dispute; rather, it involves commitment to the idea that it is uncertain whether cognitivism is appropriate to the matters in dispute. Consequently, this suggestion, even if it works for some cases, will not extend to disputes that are unproblematically cognitive, or to disputes in which the cognitive elements can be satisfactorily disentangled.

5. Religion

When we turn our attention to religious disagreement, it is clear that some of the proposals outlined in the previous section are not worthy of serious attention. In this section, I propose to examine a little more carefully the views about the epistemology of religious disagreement that seem to me to be worthy of further serious investigation. I will then close with a brief remark about the ways in which the issues for religion might be different from the issues that arise for politics.

It is not, I think, even remotely plausible to suppose that religious disagreement is primarily—or even in large part—a matter of merely verbal disagreement. Nor, I think, is it plausible to suppose that cognitive elements in religious belief cannot be satisfactorily disentangled from non-cognitive elements. However, it seems to me to be plausible to suppose that all of the following kinds of factors are present in at least some cases of religious disagreement between doxastic peers: (1) there are differences in starting points—i.e. differences in prior probabilities, doxastic presuppositions, worldviews, initial assumptions, doxastic frameworks, and the like; (2) there are differences in evidential bases for belief that are not amenable to full disclosure—of the kind envisaged by theorists such as Feldman (2007)—for the kinds of reasons mentioned earlier; (3) there is disagreement on a relatively wide range of topics—i.e. disagreement is not restricted to a relatively small number of related propositions—even though, of course, there is also wide agreement on a massive background of less relevant matters; (4) there is not total transparency of reasons for beliefs about religious matters—i.e. the reasons that one has for holding one's beliefs about questions of religion are not fully transparent to oneself, and the reasons that others have for holding the beliefs that they do on questions about religion are not fully transparent to one either; (5) there are hard questions about the extent to which the judgments that one makes about questions of religion are independent of the judgments that other people make about questions of religion; and (6) there is mutual knowledge—prior to any particular religious disagreement between doxastic peers—that there is a very wide spectrum of peer opinion on many questions about religion.

In the face of these—and other—considerations, it seems to me quite implausible to suppose that, for any given class of doxastic peers, there is at most one reasonable—rationally sustainable—position on each disputed question in religion. If there were only one reasonable—rationally sustainable—position on each peer disputed question in religion, then, plausibly, that position would involve suspension of judgment. However, it seems to me that even moderate scepticism about matters of religion based solely on grounds of peer disagreement would be a clear case of doxastic

servility. So it seems to me that the kinds of considerations listed above do push in the direction of allowing that there can be reasonable—rationally sustained—peer disagreement on questions about religion.

Of course, even if this is right, there are further questions that immediately come into view. If I grant that a doxastic peer reasonably disagrees with me on a question about religion, then what grounds do I have for supposing that I am right and my doxastic peer is wrong on that question? As we saw above, some theorists suppose that we need to say something like this: I have a special insight that is denied to my doxastic peer, and that explains why I am in the privileged position of getting things right. I think that I should say something else. I have reasons for holding the view that I do; given that I have reasons for holding the views that I do, I have reasons for supposing that my views are true (and that competing views are false). When I allow that you have reasons for holding the views that you do, I allow that you reasonably take yourself to have reasons for believing what you do. However, by my lights, while you reasonably take yourself to have reasons for believing what you do, your beliefs—insofar as they disagree with mine, and insofar as you are unable to furnish me with reason to change my mind and agree with you—are false. Of course, in the nature of the case, when I try to give you reasons to change your mind, you find the reasons that I give insufficient (and reasonably so): but that you (reasonably) don't change your mind in the face of putative reasons that I try to give to you simply doesn't give me a reason to reconsider my own views. As I see it, it is just a mistake to think that, because there is a doxastic symmetry in our position that arises from our being doxastic peers, I cannot simply rest on my first order reasons for believing as I do in order to explain why it is that I am right and you are wrong. My grounds for believing that I am right and you—my doxastic peer—are wrong are just my grounds for believing as I do.

Perhaps it might be objected that this just sounds like dogmatism: I'm right because I'm right. But that's not what I said. In order for me to have a reason to change my mind, I need to be given a reason, i.e. something that interacts in the right kind of way with the reasons that I currently have for believing as I do, and that leads me to change my view. The mere fact that you disagree with me—even if you are my doxastic peer or my doxastic superior—does not give me a reason to change my view when it comes to questions about religion.

Perhaps it might be objected that this still flouts conditions on proper explanations. Given that I recognise the doxastic symmetry of our situations, surely I need to be able to supply some further explanation of why it is that I ended up with the true belief. After all, you can certainly make the same response in justifying your refusal to change your view: your grounds for believing that you are right and that I—your doxastic peer—am wrong are just your grounds for believing as you do. However, at the very least, we need to distinguish questions here. In one sense, I can answer the question about how you ended up with false belief by telling an aetiological or historical story drawn from among the things that I believe. By my lights, I have good reasons for supposing that—at least in principle—in tracing out the aetiology of your beliefs, I would come across causes of your false beliefs that do not impugn the rationality of your holding those beliefs: *here* is where you blamelessly accepted that false testimony, etc. But, in another sense, there is no further answer to be given to the question about how you ended up with the false belief: given our presumptive

doxastic symmetry, there is no story to be told that succeeds in breaking that presumptive symmetry. (By your lights, you have good reasons for supposing that—at least in principle—in tracing out the aetiology of my beliefs, you would come across causes of my false beliefs that do not impugn my rationality in holding those beliefs: here is where I blamelessly accepted the false testimony, etc.) And, of course, this is just how you would expect matters to lie if it is true that there can be reasonable disagreements amongst doxastic peers.

Doubtless there is more to be said about the matters discussed in the previous couple of paragraphs. However, I close now with an observation about another “real” case of disagreement. In the literature, there are cases in which politics and religion are lumped in together. It is not clear that this is correct. In particular, it seems worth noting that it may be more plausible to suppose that alethic impurity has significance for political disputes than it does for disputes about religion. Whether there are supernatural beings seems to me to be an entirely cognitive question; who should govern seems to me to be a question that plausibly has a significant non-cognitive dimension. Perhaps that isn’t right; but, in any case, there is at least a *prima facie* case that there is an issue here that deserves further investigation.

References

- Aumann, R. (1976) “Agreeing to Disagree” *Annals of Statistics* 4, 6, 1236-9
- Christensen, D. (2007) “Epistemology of Disagreement: The Good News” *Philosophical Review* 116, 2, 187-217
- Christensen, D. (2009) “Disagreement as Evidence: The Epistemology of Controversy” *Philosophy Compass*
- Elga, A. (2007a) “Reflection and Disagreement” *Noûs*, 41, 478-502
- Elga, A. (2007b) “How to disagree about how to disagree” in R. Feldman and T. Warfield (eds.) *Disagreement* Oxford: Oxford University Press
- Feldman, R. (2006) “Epistemological Puzzles about Disagreement” in S. Hetherington (ed.) *Epistemology Futures* Oxford: Oxford University Press, 216-37
- Feldman, R. (2007) in L. Anthony (ed.) *Philosophers without Gods: Meditations on Atheism and the Secular Life* Oxford: Oxford University Press, 194-214
- Frances, B. (forthcoming) “The Reflective Epistemic Renegade” *Philosophy and Phenomenological Research*
- Gardner, S. and Oppy, G. (in progress) “When Bayesians Disagree”
- Genest, C. and Zidek, J. (1986) “Combining Probability Distributions: A Critique and an Annotated Bibliography” *Statistical Science* 1, 1, 114-35
- Goldman, A. (2006) “Social Epistemology” *Stanford Encyclopaedia of Philosophy*
- Goldman, A. (2007) “Epistemic Relativism and Reasonable Disagreement” in R. Feldman and T. Warfield (eds.) *Disagreement* Oxford: Oxford University Press
- Gutting, G. (1982) *Religious Belief and Religious Scepticism* Notre Dame: University of Notre Dame Press
- Harman, G. (1986) *Change in View* Cambridge MA: MIT Press
- Kelly, T. (2005) “The Epistemic Significance of Disagreement” in T. Gendler and J. Hawthorne (eds.) *Oxford Studies in Epistemology 1* Oxford: OUP, 167-96
- Kelly, T. (2009) “Peer Disagreement and Higher Order Evidence” in R. Feldman and T. Warfield (eds.) *Disagreement* Oxford: Oxford University Press

- Kitcher, P. (1990) "The Division of Cognitive Labour" *Journal of Philosophy* 87, 1, 5-22
- Lackey, J. (2009a) "What should we do when we disagree?" in T. Gendler and J. Hawthorne (eds.) *Oxford Studies in Epistemology 3* Oxford: OUP
- Lackey, J. (2009b) "A Justificationist View of Disagreement's Epistemic Significance" in A. Haddock, A. Millar, and D. Pritchard (eds.) *Social Epistemology* Oxford: OUP
- Lehrer, K. (1976) "When Rational Disagreement is Impossible" *Noûs* 10, 3, 327-32
- List, C. and Pettit, P. (2002) "Aggregating sets of judgments: an impossibility result" *Economics and Philosophy* 18, 1, 89-110
- Moffatt, M. (2007) "Reasonable Disagreement and Rational Group Inquiry" *Episteme* 4, 352-67
- Oppy, G. (2006) *Arguing about Gods* Cambridge: Cambridge University Press
- Pettit, P. (2005) "When to defer to the majority testimony—and when not" *Analysis* 66, 179-87
- Plantinga, A. (1995) "Pluralism: A Defence of Religious Exclusivism" in T. Senor (ed.) *The Rationality of Belief and the Plurality of Faith* Ithaca: Cornell University Press
- Plantinga, A. (2007) "Appendix: Two Dozen (or so) Theistic Arguments" in D. Baker (ed.) *Alvin Plantinga* New York: Cambridge University Press, 203-27
- Rosen, G. (2001) "Nominalism, Naturalism, Epistemic Relativism" *Philosophical Perspectives* 15, 69-91
- Shogenji, T. (2007) "A Conundrum in Bayesian Epistemology of Disagreement", ms, www.ric.edu/tshogenji/
- Sosa, E. (unpublished) "The Epistemology of Disagreement" <http://philpapers.org/rec/SOSTEO-2>
- van Inwagen, P. (1995) "Quam Dilecta" in T. Morris (ed.) *God and the Philosophers: The Reconciliation of Faith and Reason* Oxford: Oxford University Press
- van Inwagen, P. (1999) "It is wrong, everywhere, always, and for anyone to believe anything upon insufficient evidence" in E. Stump and M. Murray (eds.) *Philosophy of Religion: The Big Questions* London: Wiley-Blackwell, 273-85
- White, R. (2005) "Epistemic Permissiveness" *Philosophical Perspectives* 19, 445-59