

This is a preprint of an article whose final and definitive form will be published in the Australasian Journal of Philosophy; available online at <http://www.tandf.co.uk/journals>.

PEER DISAGREEMENT AND TWO PRINCIPLES OF RATIONAL BELIEF

Theodore J. Everett

This paper presents a new solution to the problem of peer disagreement that distinguishes two principles of rational belief, here called *probability* and *autonomy*. When we discover that we disagree with peers, there is one sense in which we rationally ought to suspend belief, and another in which we rationally ought to retain our original belief. In the first sense, we aim to believe what is most probably true according to our total evidence, including testimony from peers and authorities. In the second, we aim to base our beliefs only on objective evidence and argumentation, even if that lowers the probability of their being true. The first principle of rational belief tends to serve the short-term epistemic interests of individuals, while the second tends to serve the long-term epistemic interests of both individuals and groups. The best way to reconcile these principles in cases of peer disagreement is to associate them with two corresponding species of belief, here called *perception* and *opinion*.

Keywords: disagreement, epistemic peer, rationality, belief

1. Peer disagreement

My brother and I have been playing a game for many years in which we mentally compute square roots. We take a random number between 1 and 10,000, set a timer for 30 seconds, and

see who can get closer to the exact square root. So far, we seem to be about evenly matched, each of us having won about half of all the rounds we've played that didn't end in ties. So, just before each round, I believe there is about a 50% probability that I will win, and just after each round, prior to checking, I believe that there is about a 50% probability that I did win. I know how I go about computing these square roots, which is by progressive estimation: I make a first guess, take the square of that number, adjust the guess, take the new square, and so on, as many times as I can in thirty seconds. But I have no clear idea how my brother operates – he claims that he just 'sees' the numbers like the character in *Rain Man*, though he is not as good at it. Yet the fact that I have no idea how my brother's mind works seems to have no bearing on what I rationally ought to believe about who is more likely to be right. In cases like this, at least, it seems that nothing trumps basic inductive or probabilistic reasoning. Therefore, when my brother and I agree about a square root, I feel confident that my method has produced the right result. But when we disagree, as long as there is nothing special about that particular round of the game, I always suspend belief until the question is settled.

I have an old friend from graduate school with whom I play a similar game, except that instead of figuring square roots we try to figure out the concepts of knowledge, truth, and justice, and instead of thirty seconds we take thirty years. Philosophy has less definite answers than arithmetic, so it is harder for me to tell inductively which of us is more likely to be right when we find, as often happens, that we disagree. But I have plenty of indirect reasons to think that my friend and I are about evenly matched in this game. We have received similar scores in other games of mental skill like GRE exams, we have had similar records in college and in the same Ph.D. program, and we have had similar careers since then, including publishing articles in many of the same journals, and we seem to be about equally (dis)respected as philosophers by our

mutual colleagues and friends. Moreover, we have had all kinds of arguments over the years in which one of us has convinced the other that he was misinformed or making some kind of error in reasoning, and these have been about evenly balanced between the two of us. So, even in the absence of definite judgments on most of our major disagreements, I have plenty of reason to see my friend as an *epistemic peer* in matters of philosophy, that is, someone who is about as likely to be right as I am about the issues in question.¹ But when my friend and I disagree in this game, I do not suspend belief about who is right, the way I do in the square-roots game with my brother. Instead, I stick to my guns, just as my old friend sticks to his. Each of us does his best to poke dialectical holes in the other's position, and to defend his own with solid evidence and valid arguments. This is a generally pleasant, stimulating way for us to spend time together, and I suppose it satisfies a certain drive for competition in two unathletic people. But it isn't really just a game, and we're not just being stubborn for the fun of it. We also stick to our guns because we each think that we're *right* about the point in question and that the other *ought* to agree with us.

When I think about these disagreements more detachedly, though, it seems to me that I ought to believe that my probability of being right is no better than about 50%, conditional on either of us being right, just as I do when disagreeing with my brother over square roots. And in an abstract sort of way, I guess I do. That is, I know that I can't reasonably say that I am a better philosopher than he is overall, or that I'm right and he is wrong in every disagreement that we have. It seems to follow that I ought to believe that when the truth is finally revealed, it is more

¹ Here I am following Elga [2007] and Enoch [2010] as opposed to most of the current literature, where epistemic peerage is defined to include not just equal overall reliability, but also possession of the same evidence and the same epistemic virtues.

or less a coin toss which of us is going to be proven right, supposing either of us is. But somehow that is *not* what I believe – not how I feel, speak, or behave – in each case as it comes along. In any particular disagreement, I believe that I am right and he is wrong. Not only that; I also usually feel that he is being rather dense about the point in question, and that while I understand his arguments, he isn't really understanding mine. He lets me know that he feels the same way about me, so we are clearly still in a symmetrical relation. But this doesn't seem to matter to my confidence that I am right and that it is reasonable for me not to give in. So, I hold on to my position and keep arguing the point, just as he does, until we run out of time, or steam, or patience, or (occasionally) come to agreement on the issue.

What accounts for the difference in my attitude towards these two sorts of disagreement? I accept that I am often wrong in my mathematical computations, and more or less cheerfully suspend belief in peer disagreements about arithmetic. But I insist on retaining belief in the analogous peer disagreements in philosophy, even though I know that I am also likely to be wrong about philosophy – indeed, more likely in philosophy than in arithmetic. Why should I think this way? It can't be just because my philosophical reasoning is transparent to me while my friend's is partly opaque, because the same is true in the square-roots game with my brother. For I have no conception at all of how my brother makes his calculations; I just know that they turn out to be about as reliable as mine. And it can't be just because I often think I see what is wrong with my friend's positions, since he also often thinks he sees what is wrong with my positions, and all of my evidence, including our long history of partly-conclusive arguments, shows that my guesses about what is wrong with his arguments are no more reliable than his guesses about what is wrong with mine. Even when something seems perfectly obvious to me, experience has shown that this is not a better guide to the truth than what seems obvious to him.

So, why should I keep sticking with my own beliefs in each of our disagreements – and thinking that this is the right thing to do – instead of humbly accepting that my friend’s beliefs are just as likely to be true as mine are?² To believe a proposition surely entails believing that it is at least probably true. How can I do so rationally when my total evidence strongly implies that it is just as likely to be false?

Here is the basic problem of peer disagreement as I see it, in the form of a loose paradox:

(1) When you are likely to be wrong, you ought to suspend belief.³

² For arguments that we are rationally bound to give a disagreeing peer's position equal weight with our own, see Feldman [2006] and Elga [2007]. For an argument that we are rationally bound to ignore such evidence, see Kelly [2005]. For arguments that it should be considered but not given equal probabilistic weight with our own prior beliefs, see Lackey [2008], Kelly [2010] and Enoch [2010]. For an argument that it should be given *almost* equal weight, see Cohen [2013]. Here I am taking the intuitive appeal of what is generally called the Equal Weight View for granted (on the model of simple disagreements in arithmetic) in the interests of developing a new, broader analysis of peer disagreement, not pretending to have refuted these or any other well-elaborated positions. For what it is worth, it seems to me that there are two good reasons for departing somewhat from the Equal Weight View: first, that we know our own immediate state of mind better than that of our peer opponent, so we should discount his testimony by the (typically small) differential probability that his is lying or somehow impaired [Lackey 2008]; and second, that rational equilibrium requires us to reduce his prior status as a peer (typically only slightly) once we discover that he disagrees with our prior position on the issue in question [Kelly 2010]. In any case, my argument here does not depend on weighting the views of peers absolutely equally with our own. A weaker principle would do, to the effect that a peer's disagreement carries *some* epistemic weight all by itself, regardless of whether it seems right or makes sense to the believer.

³ For simplicity's sake I am leaving this vague. There is no single general threshold of likelihood of error that ought to trigger categorical suspension of belief, though clearly, the more likely you are to be wrong, the more reason you have to suspend belief. On the roughly Bayesian sort of approach I actually favour, our degrees of confidence

- (2) When peers disagree with you, you are likely to be wrong.
- (3) (Therefore) When peers disagree with you, you ought to suspend belief.
- (4) (But) When peers disagree with you, you ought *not* to suspend belief.

The argument from (1) and (2) to (3) is plainly valid. Moreover, when we think about simple matters like arithmetic, or in general terms about our own fallibility, even within our range of expertise, the argument seems perfectly sound. So, it looks like (3) is true. But when we think about particular disagreements in philosophy, (3) seems to be false and (4) seems to be true instead. What is going on? I think the problem is that (3) is always true in a sense, but (4) is sometimes true in another sense, where the ambiguity hinges on different principles that govern how we rationally ought to think, or, to put it linguistically, on different uses of the epistemic 'ought'.

We all accept that the word 'ought' can be used in different ways depending on the purposes that speakers have in mind. Sometimes we mean what people ought to do in order to be morally good or to have done the morally right thing – what is called the *moral* 'ought'. But there are many other uses of the word 'ought' that depend on other goals that speakers might have in mind. Thus, you might tell people that they ought to stop smoking, meaning not that this will make them morally better people, but just that they'll be better off in terms of their own interests – the so-called *prudential* 'ought'.⁴ This makes it easy to equivocate sometimes, for example as

should be adjusted continuously according to our total evidence at any moment, so there is no sharp distinction between suspending and not suspending belief.

⁴ Other senses of the word 'ought' have no special label, but they are just as useful. In fact, any sort of goal or interest can define its own sense of the word 'ought'. For a few examples, we might say that hot cocoa ought to have a little cinnamon in it, meaning that it tastes better that way; that a hockey puck ought to be three inches in diameter,

to whether a soldier facing battle ought to run away or ought to stand and fight. Here is another little paradox, analogous to the one above about peer disagreement:

- (1a) When you are likely to be killed, you ought to run away.
- (2a) When you are attacked in battle, you are likely to be killed.
- (3a) (Therefore) When you are attacked in battle, you ought to run away.
- (4a) (But) When you are attacked in battle, you ought *not* to run away.

This is not much of a paradox, of course, just an equivocation in (3a) and (4a) between the prudential and moral senses of 'ought'. In the prudential sense of 'ought', where the implicit goal is serving one's own interests, a person in great danger generally ought to run away if he can. But in the moral sense of 'ought', where the implicit goal is serving the general interest of his country, a soldier ought in many of the same cases to stand and fight, even at great risk to his own life.⁵

Here is the problem with this analogy. What is at work in the paradox about peer disagreement is neither the moral nor the prudential 'ought', but what we call the *epistemic* 'ought', meaning what people ought rationally to believe, given the goal of knowledge, or at least well-justified belief. The immediate problem, then, is that the 'ought' seems to be unequivocally epistemic in both (3) and (4), so the analogy seems not to help at all. When we think we ought to

meaning that this meets official standards; or that Keyser Söse ought to murder his own family, meaning that this makes *The Usual Suspects* a more interesting movie.

⁵ In case this isn't obvious, I do not mean to imply that there are no moral arguments for soldiers to run away or prudential arguments for them to stand their ground. I am just assuming a traditional view of such things in order to make my point.

suspend belief in arithmetic disagreements, we mean simply that this is the rational thing to do. And by the same probabilistic reasoning, it seems that we ought to suspend belief about our controversial philosophical positions as well, since it makes no sense to say that we believe something without believing that it is at least probably true. But when we think we ought to retain belief in particular philosophical disagreements, we also mean that this is the rational thing to do, despite the implication that we therefore believe we are probably right.⁶ That is, we do not think that we're sticking to our guns in arguments just because it makes us feel good, or because we don't want to be seen as wishy-washy, or anything of the sort. We believe that we are being perfectly rational, and we explain our stands with rational arguments, not with appeals to other values. So, we are left with the problem of reconciling our confidence in these beliefs with the evident fact that they are likely to be wrong in cases of peer disagreement.

2. Two principles of rational belief

The problem seems to turn on which of these two attitudes, humility or self-assurance, we believe that rationality demands in cases of peer disagreement. Philosophers seem to be about evenly split on this question, with some arguing that rationality requires suspension of belief (or 'conciliation' as much of the literature has it) in all such cases, and others arguing that sticking to our guns (or 'dogmatism') is what is rational at least some of the time. But should we try to make a single sort of judgment as to what is rational in cases like this? Is rationality even a univocal concept? In general, it seems to mean something like this: reasoning in a way that leads reliably

⁶ Van Inwagen [2010: 28] goes so far as to claim that he *incapable* of accepting that his confidence in being right in his disputes with peers is irrational.

to true beliefs. But which true beliefs are we talking about, exactly? True beliefs for whom, and when, and under what conditions? Different answers to these questions could yield many conflicting judgments as to what is rational in this or that particular case. I believe that there are two philosophically essential ways to answer these questions, and that the often but not always inconsequential difference between them accounts for our conflicting intuitions about peer disagreement. So, let me propose two principles of rational belief, each of which I think defines one sense or sub-sense of the epistemic 'ought'.

The *principle of probability* is that you ought to believe whatever is most likely to be true, given your total pool of evidence. More precisely, you should believe with greater confidence whatever is more likely to be true, given the total evidence available to you, and to adjust that confidence accordingly whenever new evidence appears.

The *principle of autonomy* is that you ought to base your beliefs (or degrees of belief) solely on objective evidence, using your own best reasoning and judgment. You should consider the arguments of others on their merits, but you should not allow the simple probability that they are right to influence your thinking.

These principles determine different epistemic 'oughts' because they reflect different fundamental epistemic interests. The principle of probability reflects the goal of maximally justified belief at any moment, which is primarily a goal of individuals who need to act. When thinking only of your own immediate probability of being right on any issue, you should consider all of the evidence available to you, including testimony from a peer or any other source

that you have reason to consider somewhat reliable. There is no reason to rule out any evidence at all, if the only thing you care about is the most probable truth *right now*. So, if you are being forced to bet your life, say, on some unestablished fact, then you should typically weigh the testimony of your epistemic peers more or less equally with your own prior opinion on the matter, and you should weigh more heavily the testimony of your epistemic betters, if you have any, even in your own areas of expertise. Thus, in medical decision-making where lives may be at stake, doctors are ordinarily expected to follow protocols or 'standards of care' that the consensus of their peers say yield the highest probability of good results, rather than their own, perhaps eccentric, theories.⁷

But we have other epistemic goals as well. We do not just want to place bets on which existing ideas are most probably right, but also to produce, defend, and criticize ideas and arguments in ways that ultimately lead to greater knowledge for ourselves and our societies. When faced with intellectual problems and puzzles, we try to *solve* them, not just to guess at what theories will turn out to be true. There are two connected reasons for this. One reason is that we desire as individuals to understand the world, not just to play the market, as it were, of probabilities. There is no knowledge worthy of the name without at least a fair degree of understanding. For example, I can say that I believe in quantum mechanics because physicists tell me that this is the best-established theory in their field, and I have reason to suppose that they are probably telling the truth. But I have only the wispiest notions of wave-particle duality and other concepts integral to quantum mechanics, hardly enough to say that I believe anything about quantum mechanics itself, as opposed to just believing that there is a theory *called* 'quantum

⁷ Even television's Doctor House typically allows his staff to follow protocol until all of the likely diagnoses have been tested and failed, at which point he intervenes and saves the day with brilliant hunches.

mechanics' that is probably true. If I have any real interest in physics, such degenerate beliefs are of essentially no use to me. Even where I can clearly grasp the major claims involved (as with the thesis of anthropogenic global warming, say), I still can't claim to understand the issue as a whole, let alone know whether the thesis is true, without examining the arguments objectively. (For me even to *know what I am talking about*, as it is commonly put, means at a minimum that my statements must make sense to me along objective lines of reasoning.) Whether my peers agree or disagree with me has little bearing on the matter, except as it gets me to notice their ideas and arguments, which I can then evaluate strictly according to their merits. To the extent that rational belief aims at real knowledge, then, as opposed to mere successful bets on propositions, the principle of autonomy would seem to trump the principle of probability.

Our other essential reason for thinking autonomously is, ironically, that our deepest intellectual problems are typically too subtle and complex for one person to solve.⁸ We must think for ourselves in order that as many plausible theories as possible can be criticized and tested by other thinkers, also acting independently, in the expectation that the truth will someday emerge from this collective competition. No doubt, some philosophers or scientists are better at constructing theories than others, in the sense of being more likely to be proven right over the long run. There might even be some one philosopher superior to all the rest of us, so that if we had to bet serious money on any particular theory being true, it would be rational for us to bet on that person's theory rather than one of our own. But why should we think we have to make such bets? We are not in this business just to gamble on which theories will turn out to be right. We are in it to work on solving hard problems over a long time, both as individual philosophers and collectively, as members of the philosophical profession. It would be absurd for us to leave the

⁸ Elgin [2010], citing Kitcher [1990], raises this point in defence of retaining belief in cases of peer disagreement.

whole business to a single most-probably-correct philosopher, because even the best of us makes plenty of mistakes, and even the least of us is capable of contributing useful ideas to the ongoing discussion. For the same reason, we should not be discouraged when it turns out that our epistemic peers disagree with us. *Of course* they do, because it is part of our very job to come up with new ideas and new objective arguments to back them up. As philosophers, we are producers and critics of ideas, not just consumers, so if we do not think for ourselves, then we are not being responsible, effective members of our community.

Much of this creative sort of thinking can in fact be done in an entirely hypothetical spirit, with no violation of the principle of probability. In working on difficult problems we can, and often do, experiment with theories we consider unlikely to be true and see what develops, in the confidence that peers are working on other (and perhaps more plausible) conjectures. There is no reason in principle that we should believe in any of these theories to a degree beyond what all the evidence, including testimony from our peers, entails. In fact, if we are all completely rational and fully informed of each other's evidence and reasoning, we ought ideally to be able to agree on a single, shared subjective likelihood for each hypothesis that we consider, and this would not prevent us from continuing to work towards a more permanent consensus. Much current scientific practice is already like this, more or less. In industry, for example, the scientist is someone with a job; his leaders give him a project to work on, and whether he personally thinks the project will succeed is hardly relevant to what he has to do. And in medicine, researchers are particularly conscious of the complex social nature of their work, accepting that unlikely possibilities need to be carefully ruled out for the sake of completeness in their broad-based investigations.

In philosophy, though, and in more revolutionary science, we have four strong reasons to adhere to the principle of autonomy, working our own theories out as individuals regardless of the level of agreement from our peers. First, though theoretical diversity can in principle be maintained by people working with hypotheses they don't believe, philosophers are not just motivated to be helpful in communal projects; we also seek for truth and understanding for ourselves. So, it is natural for us focus our attention on hypotheses that strike us independently as the most probably true, rather than work against our own epistemic interests on theories we consider less likely to pan out. Second, as philosophers we are expected not just to produce new theories, but also to promote them and defend them in the public arguments that constitute our testing system. We are poor actors, most of us, so a good measure of sincere belief is usually needed for us to be effective advocates, especially for complex theories that demand years of debate. Third, we are also philosophically more competent defending our own theories sincerely than our opponents' hypothetically, because our own theories articulate perceptions of the way things really are, while our opponents' typically appear to us as sets of propositions that are false at best, and that at worst don't even make sense. And fourth, to develop and promote dissenting theories in particular demands persistence in the face of not just widespread disagreement, but often also ridicule, rejection, and even persecution from our peers, as witness Socrates or Galileo, and this is almost impossible absent the conviction that we are at least probably right.⁹ Not as a matter of ideal epistemology, then, perhaps, but psychologically, at least, it seems that we must believe in what we say in order to say it maximally well, to persist in its autonomous

⁹ Interestingly, prior to his condemnation Galileo had long been encouraged by Church authorities to pursue his heterodox researches in a hypothetical vein, subject to ultimate approval by his friend the Pope. But he could never bring himself to say that he was only 'saving the appearances' with his model of the universe. What propelled him in his work, and what ultimately got him into so much trouble, was that he *believed* in it.

development over a long career, and to withstand the consequences of upsetting other interested parties.

Here, then, is my preliminary solution to the problem of peer disagreement. We have two different, equally useful principles that govern rational belief formation, and these define two corresponding uses of the epistemic 'ought'. Both employ the same inductive and deductive methods, so there is no difference in the rationality *per se* of these two principles; the only essential difference lies in what gets counted as appropriate evidential 'input' to the rational machinery.¹⁰ One principle takes in all available evidence, including testimony from reliable sources, and produces probabilized bets on facts. The other excludes evidence derived solely from testimony, and produces arguments and theories necessary for both understanding and objective progress in philosophy and much of science. *Qua* mere consumers of theories, then, we ought to suspend belief on probabilistic grounds when confronted with disagreement from people as likely as ourselves to turn out to be right. *Qua* producers and defenders of theories and *qua* seekers of understanding, we ought to stand by our own beliefs until we are convinced to yield them on objective grounds.

3. What to believe

There is something unsatisfactory about a theory that posits two contradictory things that rationally ought to be believed, without saying which is finally to be preferred. In cases of peer

¹⁰ My position is intended to be consistent with the 'Uniqueness Thesis' that there is only one rational conclusion to draw from any set of evidence. But this depends on what we take to be the pool of evidence that matters. On my view, we can in fact rationally draw different conclusions from a given pool of evidence, but only in the sense that it is sometimes (i.e. under the principle of autonomy) epistemically proper to ignore some of that evidence.

disagreement, even allowing that opposed beliefs are rational in different ways (or that we ought to believe them in two different senses of the word 'ought'), the question remains: what, after all, on balance, should we *really believe*? It looks like the original problem of peer disagreement must now be recapitulated, not directly in terms of which belief is rational according to common sense (for they both are), but indirectly, in terms of which of standard of rational belief takes precedence when they conflict. Taking for granted that we cannot rationally hold two contradictory beliefs at once, there are four possible coherent answers to this latter question: probability, autonomy, neither, and both.

The first option is to say that in cases of peer disagreement, the only rational thing for us to do is to follow the principle of probability inwardly, respecting what we really believe, even while following the principle of autonomy outwardly in our debates with peers. On this approach, we ought to view our independently-developed arguments and theories with sceptical detachment, accepting that we are likely to be wrong while continuing to work on making concrete sense of the matter for ourselves and others. So, the achievements of Einstein and Wittgenstein are great, and may well have depended causally on an autonomous approach; nevertheless, these thinkers had no rational warrant to believe that they were right. So, in an epistemic if not in a practical or moral sense, they ought not to have believed in their own theories.

The second approach is to say that our real beliefs are our autonomous beliefs, and that mere probabilized statements that derive from testimony ought not to count. We will have practical reasons for considering peer disagreement when we need to guess at the truth for purposes of action. In philosophical disagreements, though, where no real decisions are required, the principle of probability has no force at all. We will observe that other people just as

sharp and well-informed as we are see things differently, but we will have no strictly epistemic need to reconcile their different views with ours in terms of probabilities. Odd as it sounds, then, we can in fact believe something without believing that it is probably true.

The third possible solution is to claim that there is no fact of the matter as to which principle is more important, so that neither of the first two approaches is correct. Instead, belief can be determined in cases of peer disagreement only by the interests of the believer. If the believer seeks to be most-probably right, then he should follow the principle of probability. If he wants to understand things and develop new ideas, then he should favour the principle of autonomy. Our soldier in battle ought prudentially to run away, and ought morally to stand and fight – but how can it be clear which one he ought to do, all things considered? Unless there is good overall prudential reason for him to prefer the moral action, or good moral reason to prefer the prudent one, the soldier seems to be left with a brute choice to make, not a rational decision.¹¹ The same can be said to hold for people like ourselves in cases of peer disagreement: there is no other choice but just to *choose* what we believe.

The fourth way around the problem is to claim that both principles can safely be followed at the same time, because they never actually produce contradictory beliefs. In fact, thinking autonomously will always maximize the probable truth of our beliefs. It is hard to see how this thesis can make sense as a general rule, for it seems to imply that each of any pair of disagreeing peers is more probably right than his opponent.¹² But each of us could separately work around

¹¹ Some philosophers do argue that rationality itself favours one or the other preference, say Kant for morality and Nietzsche for a kind of prudence. I am just pointing out the way that things intuitively seem.

¹² It is, of course, rationally possible for peers to have different *subjective* probabilities for the same proposition, at least initially. The intuitive problem is that each of them should also be able to view their disagreement from a shared, *objective* point of view, and this would seem to wipe out each of their preferences for their own prior beliefs,

this implication by denying that we have any peers at all who disagree with us. If we claim that we can follow both principles together and end up with consistent beliefs, then we must accept that the mere fact that others disagree with us excludes them categorically as epistemic peers.

None of these options strikes me as satisfactory. The first approach fails in privileging the probabilist betting-on-things-right-now conception of rational belief over the constructive sort of rationality required both for understanding and for major progress in philosophy and science. This makes good blackjack players rational and great thinkers like Galileo not, which is a hard consequence to swallow intuitively. If philosophers and scientists aren't being rational in thinking for themselves, we need another word of epistemic praise that's just as good. The second approach has the opposite problem: it may be rational for me to stick to my guns in philosophical disputes, but it is surely still *irrational* for me to do so in peer disagreements over things like arithmetic, where comparative track-record constitutes most of our evidence. The third approach allows us to follow both principles, which is good, but forces us to choose what to believe whenever they conflict, according to our interests. What if our interests lie primarily in having rational beliefs? The third approach permits no answer. It also joins the second approach in licensing belief in things that we do not believe are even probably true. And the fourth approach, to say that nobody who disagrees with us counts as an epistemic peer, cannot succeed because few serious philosophers are quite so arrogant; and anyway, it's not a general solution. Many are working on the problem of peer disagreement, and this approach could only satisfy one person at a time.¹³

on pain of contradiction. Enoch [2010] argues that the first-person point of view is nevertheless dominant for purposes of rational belief revision.

¹³ Some philosophers do appear to think and act this way, as if they believed that they were literally peerless in their areas of interest, though few are so ungracious as to admit it. Among those decently troubled by their tendency to

4. Two ways of believing

The only possible alternative approach to these four is to give up our assumption that beliefs must be consistent to be rational. This looks like nonsense on its face if we think of belief in the normal way, that is, univocally, so that for every proposition you consider, either you believe it or you don't. But let me try to distinguish between two sorts of belief or ways of believing, one for each of the two rational principles to govern exclusively. If a theory of this sort could be developed in a reasonable way, both principles could then be followed in peer disagreements without yielding contradictory beliefs of either of the two types.

There is evidence for just the right sort of ambiguity in the way that we distinguish what we call *opinions* from other beliefs. When we find ourselves challenged by peers who disagree with something we have said, we often retreat to some extent by saying things like, 'Well, I was only stating an opinion.' This suggests that the beliefs we call opinions leave room for debate and doubt in a way that other beliefs do not. In fact, unless we are acknowledging or anticipating disagreement of some sort, it seems to me we never call a belief an opinion. So we might think of something like my belief that Canada will someday rule the world as just a belief, not an opinion, until we discover that there are peers of ours who disagree. At that point, the belief in question either survives as an opinion or it stops being a belief at all – depending on whether or not we desire to maintain it in the face of disagreement from our peers. Let me extend the ordinary meaning of the word 'opinion' to include such potential as well as actual cases, so that

take the fourth approach while not really believing that they are smarter than everybody else is Fumerton [2010: 103-105], who wonders whether his well-known tenacity in philosophical discussions makes him seem like a 'jerk' or 'egomaniac'. It does not, but he is right to be puzzled by this fact.

any belief that you would tend to retain in the face of disagreement will count as an opinion in the sense I mean. And let me say that to believe something in this sense is to *hold it as an opinion*, or simply to *hold it*, so as to avoid the clumsy word 'opine'. So, I will say that I hold that Canada will someday rule the world, regardless of whether anybody disagrees. All beliefs are subject to disagreement, of course, but it is only beliefs we have worked out for ourselves to some extent that we are liable to maintain when faced with disagreeing peers; otherwise, we would have no concrete arguments to make. Opinions in my extended sense may be conceived, then, simply as beliefs derived according to the principle of autonomy.

Another sense of the word 'belief' respects the way that we perceive the world after all evidence has been considered. As I have said, beliefs in this sense are sometimes only probabilized repetitions of things that we have been told, with little understanding or autonomous justification. To reprise my earlier example: I believe that there is such a thing as wave-particle duality, based only on testimony from experts. If a peer were to challenge this belief of mine, I could hardly retain it as an opinion since I have no independent grounds at all for arguing the point. All I could reasonably do is lower the subjective probability I assign to the statement that wave-particle duality exists in light of the contrary testimony from a peer. Nevertheless, in the absence of actual peer contradiction, I believe that wave-particle duality exists, just in that it forms a part, however poorly integrated, of my probabilized model of the world. So, if I had to bet for or against the proposition that there is such a thing as wave-particle duality, I would bet for it – which is really all that believing something in this way amounts to. For want of a better single word, let me label all such beliefs *perceptions*.

Perceptions and opinions are best understood things of the same intrinsic type, differing proximately in their being derived from unrestricted and restricted sets of evidence, respectively,

and ultimately on their serving different epistemic functions.¹⁴ They should not be seen as mutually exclusive classes. In fact, in most cases most of the time, there is no concrete difference at all between what we perceive and what we hold as an opinion. The two ways of believing only tend to come apart under the stress of peer disagreement, when we need to separate the probable from the productive and well-understood. Otherwise, beliefs are just beliefs.

For clarity's sake, here is an outline of a theory about peer disagreement that connects this distinction between types of belief to our original distinction between principles of rational belief.

¹⁴ Elgin [2010], following Cohen [1992], distinguishes for different reasons between 'belief' as an involuntary *feeling* that something is so, and 'acceptance' as a voluntary, action-guiding sort of assent. She argues that the problem of rational belief in cases of peer disagreement is really about acceptance rather than what she calls belief, since a claim that we *ought* to believe something implies that we *can* believe it, and we cannot always change our beliefs when we discover that they are irrational. It seems to me that the purely epistemic 'ought' does in fact always imply the epistemic 'can' – just not the *psychological* 'can'. For analogy, the fact that someone ought to castle at a certain point during a game of chess clearly implies that he can castle, that is, consistently with the rules of chess. Perhaps he oughtn't to castle as a moral or prudential matter (say, his family's lives are hostage to his losing the game), and perhaps he cannot physically move his king and rook (say, he is incapacitated by a stroke). But he still *can* castle in the same restricted, formal realm in which he *ought* to castle, i.e. the game *qua* game. Similarly, if a person ought to believe something (say, that his spouse is cheating on him) in the restricted, purely epistemic sense that it is rationally necessary for him to believe it in light of the available evidence, this implies that he can believe it in the corresponding sense that it is rationally possible for him to believe it, but it does not imply that he either can or ought to believe it in the practical world.

(1) There are two principles of rational belief: the principle of probability and the principle of autonomy.

(2) In ordinary reasoning, these two principles produce the same beliefs.

(3) In cases of peer disagreement, however, the two principles tend to conflict, producing what seem to be contradictory beliefs.

(4) But there are also two kinds of beliefs: perceptions, which are rationally governed by the principle of probability, and opinions, which are rationally governed by the principle of autonomy.

(5) In ordinary reasoning, there is only one relevant kind of belief, since the two governing principles produce the same results. That is, perceptions and opinions are ordinarily the same things, which we simply call beliefs.

(6) In cases of peer disagreement, however, the two sorts of belief come apart. We perceive one thing according to the principle of probability, and we hold something else according to the principle of autonomy.

(7) In such cases there is no one thing that we believe *simpliciter*. Rather, we believe two different things in different senses of the word 'believe'.

(8) There is no reason to favour one principle or one sort of belief over the other as uniquely rational in situations of peer disagreement.

I believe that this points to a satisfactory solution to the problem of reconciling the two principles of rational belief in situations of peer disagreement, hence to the problem of peer disagreement itself. It should be obvious that it works technically, that is, that it provides at least a superficially coherent way of structuring the necessary concepts. It also seems to me that it can serve to validate both of our contrary intuitions in peer disagreements without doing too much damage to the ordinary meanings of the words involved, and without forcing unnecessary choices about what to believe. It allows us to make sense of why we tend to suspend belief when faced with disagreement in matters like the accuracy of square root calculations, since there is nothing at stake for us here in the relevant sense: we have no reason to persist in working out or testing or asserting a personal position on which number is correct, hence no actual opinion on the matter. At the same time, we can understand why we tend to stick to our guns in things like philosophical disagreements, for here we do have a personal epistemic interest in constructing an integrated understanding of the point in question, and a further social interest in contributing autonomous ideas and arguments to the more general discussion. Thus, we can acknowledge in an abstract way that we are likely to be wrong, all things considered, while still sincerely urging that our own opinion is correct. As to which of these things we believe, the answer is that we believe both things in different ways. As to which one we *really* believe, there is no answer and no need for one.

Note that this theory resolves the problem with saying that we can believe what I am now calling opinions despite their likelihood of being wrong, namely that by definition, to believe

something entails believing that it is at least probably true. In this theory, both types of belief conform to that principle: if we perceive something then we perceive that it is probably true, and if we hold something then we hold that it is probably true. The problematic inference from believing (in the sense of holding) something to believing (in the sense of perceiving) that it is probably true is now blocked.

The theory also suggests a neat solution to the problem of believing that we are right in each case of peer disagreement separately, but wrong in many of these cases when they are viewed as a group. Considering our disagreements with our peers in general, we tend to agree that we are just as likely to be wrong as they are, for this is a simple statistical perception governed by the principle of probability. At the same time, though, when each particular disagreement occurs we tend to insist that we are right about the point in question, for these are matters of opinion governed by the principle of autonomy. Both attitudes are rational; no paradox results as long as we discriminate between the two types of belief.

There is an intuitive cost to all of this, of course, in that it makes the concept of belief equivocal, permitting more than one doxastic attitude towards a single proposition. Do we really want to say that someone can both believe and not believe the same thing at the same time *at all*, let alone rationally? On balance, yes, I think we do. But does this not entail abandoning the epistemic principle of non-contradiction? I don't think so. If our beliefs do form two overlapping sets or systems, and neither one contains internal inconsistencies, then this should suffice to satisfy the principle. It is, admittedly, intuitively odd to say that each of us has two sets of beliefs that sometimes conflict. But if we fully appreciate the different functions of perceptions and opinions in our individual and social epistemic lives, and if we value a clean-cut solution to the problem of peer disagreement, then we can probably accept this as a price worth

paying. This is, at any rate, what I believe at present. If I discover that peers disagree with me on this, then I'll perceive that I am likely to be wrong – but I will need to be shown *how* I am wrong before I give up my opinion.¹⁵

State University of New York at Geneseo

References

- Cohen, S. 2013. A Defense of the (Almost) Equal Weight View, in *The Epistemology of Disagreement: New Essays*, ed. D. Christensen and J. Lackey, Oxford: Oxford University Press: 98-120.
- Elga, A. 2007. Reflection and Disagreement, *Nous* 41/3: 478-502.
- Elgin, C. 2010. Persistent Disagreement, in *Disagreement*, ed. R. Feldman and T. Warfield, Oxford: Oxford University Press: 53-68.
- Enoch, D. 2010. Not Just a Truthometer: Taking Oneself Seriously (but not Too Seriously) in Cases of Peer Disagreement, *Mind* 119/476: 953-997.
- Feldman, R. 2006. Epistemological Puzzles about Disagreement, in *Epistemology Futures*, ed. S. Hetherington, Oxford: Oxford University Press: 216-36.
- Fumerton, R. 2010. You Can't Trust a Philosopher, in *Disagreement*, ed. R. Feldman and T. Warfield, Oxford: Oxford University Press: 91-110.
- Kelly, T. 2005. The Epistemic Significance of Disagreement, in *Oxford Studies in Epistemology, Volume 1*, ed. J. Hawthorne and T. Gendler Szabo, Oxford: Oxford University Press: 167-196.
- Kelly, T. 2010. Peer Disagreement and Higher Order Evidence, in *Disagreement*, ed. R. Feldman and T. Warfield, Oxford: Oxford University Press: 111-174.

¹⁵ Thanks for helpful comments to David James Barnett, Alan Sidelle, audiences at SUNY Geneseo and the Creighton Club, and two anonymous reviewers for the *Australasian Journal of Philosophy*.

Kitcher, P. 1990. The Division of Epistemic Labor, *Journal of Philosophy* 87: 5-22.

Lackey, J. 2008. What Should We Do When We Disagree?, in *Oxford Studies in Epistemology, Volume 3*, ed. J. Hawthorne and T. Gendler Szabo, Oxford: Oxford University Press: 274-293.

Van Inwagen, P. 2010. We're right. They're wrong, in *Disagreement*, ed. R. Feldman and T. Warfield, Oxford: Oxford University Press: 10-28.