

## Epistemic Disagreement

Suppose we mean by your epistemic peer someone who not only possesses all the intellectual virtues that you do (intelligence, perspicacity, honesty, thoroughness, conscientiousness, thoughtfulness, freedom from bias, sincerity in the pursuit of truth, etc.) but also someone who has been exposed to just the same arguments and evidence that you have.<sup>1</sup> Tom Kelly then says, “The question at issue, then, is whether known disagreement with those who are one’s epistemic peers in this sense must inevitably undermine the rationality of maintaining one’s own views.”<sup>2</sup> I think he means that the question is whether known disagreement with those *known to be* one’s epistemic peers must undermine the rationality of your own view since the mere fact that I disagree with Fred, whom I do not realize is my epistemic peer, need not undermine the rationality of maintaining my views. And as David Christensen has pointed out, the real question is whether known disagreement with known epistemic peers rationally requires you to lower your level of confidence in your relevant beliefs.<sup>3</sup> It might require that even if it does not completely undermine the rationality of maintaining those beliefs, that is, does not make it irrational to still hold them. A lowered level of confidence is compatible with your still rationally maintaining your beliefs. The level of confidence in your belief that P initially might be 0.9 and when you learn that your epistemic peers disagree, you might be rationally required to lower it to 0.75, a level that is compatible with your still rationally

---

<sup>1</sup>Tom Kelly defines “epistemic peer” this way in his “The Epistemic Significance of Disagreement,” in *Oxford Studies in Epistemology* (Oxford: Oxford University Press, 2005), Tamar Szabo Gendler and John Hawthorne, eds, pp. 174-75 and p. 168. He distinguishes his broad sense of being an epistemic peer from Gary Gutting’s narrower sense that only considers equality in epistemic virtue, not equality in the possession of evidence. See p. 168 and note 2 on p. 168.

<sup>2</sup>Kelly (2005), p. 175.

<sup>3</sup>David Christensen, “Epistemology of Disagreement: The Good News,” *The Philosophical Review*, (2007) 116 (2), pp. 187-217. On pp. 188 and 199 Christensen writes about “revising ones confidence in P” and “moving one’s belief” towards the other’s belief, respectively, rather than all-or-nothing belief, or suspension of belief.

believing that P.

Kelly argues that known disagreement with known epistemic peers need not undermine the rationality of your maintaining your belief on which your peers disagree. Why? Because, he says, your judgment that your epistemic peers have misjudged the probative force of the evidence for P (say, that Truman's desire to intimidate the Soviet Union was a motivation for his dropping the atomic bomb on Japan) breaks the symmetry between you and your peers regarding epistemic virtue and evidence possessed.<sup>4</sup> But surely that judgment does not break the symmetry if it is not justified, if you just happen to hold it for no good reason. Suppose there is evidence, e1, for P and also your evidence, e2, for your judgment that your peers have on this occasion misjudged the probative force of e1. Now if the others do not possess e2, they are not your epistemic peers in Kelly's broad sense since they are not your evidential equals. Then, of course, you might be perfectly rational in believing P and they perfectly rational in believing not-P since they may have no reason to think that they have misjudged the evidence, e1. But if they also possess e2 and you have evidence, e3, that they are your epistemic peers and that they disagree and do not believe P, why isn't the rationality of your belief that P undermined? Shouldn't you think, "I have no more reason to think that they have misjudged the evidence than that I have. I would have such reason if I had reason to think that they are not my evidential equals, that is, that they lack evidence I possess, or lack some of my epistemic virtues. But since they are my epistemic peers, I have no such reason. I thought e2 was a good reason to think that *they* misjudged the evidence, but e2 hasn't convinced them. And I *know* they are my epistemic peers. I better lower my confidence that P is true."

Kelly argues, "On the present view [his view], the rationality of the parties engaged in

---

<sup>4</sup>Kelly (2005), p. 179.

such a dispute will typically depend on who has *in fact* correctly evaluated the available evidence and who has not.<sup>5</sup> Does he mean the evidence,  $e_1$ , for P or the totality of evidence,  $(e_1 + e_2 + e_3)$ ? Christensen's example of two mathematical peers at a restaurant refutes the idea that the rationality of maintaining your belief in light of known disagreement with a mathematical peer is simply a function of who has correctly evaluated the evidence for the relevant belief.<sup>6</sup> In one of his examples A and B calculate in their heads what each person at the table owes by taking the total bill, adding a 20% tip, and then dividing by the number of people at the table. Assume that each knows that the other has a good track record with this sort of task but that they disagree on the answer on this particular occasion. Assume, also, that one of them has done the calculation correctly and so has "correctly evaluated the available evidence" in reaching his conclusion. Still, in this circumstance where that person knows that a known mathematical peer has arrived at a different answer, the rationality of his maintaining his view about what each person owes has been undermined. So the rationality of maintaining your views in circumstances where you know that known epistemic peers disagree cannot be solely a function of the evidence for your views, that is, of  $e_1$ .

Perhaps that rationality is a function of the total evidence  $(e_1 + e_2 + e_3)$ , not just of  $e_1$ . But in the context of known disagreement with known epistemic peers what evidence could there be that "on this particular occasion I have done a better job with respect to weighing the evidence and competing considerations than they have"?<sup>7</sup> It is true that a morally virtuous person need not always do the right thing and that an epistemically virtuous person need not always weigh the evidence correctly, but they must generally do that. So why think they, rather than you, are

---

<sup>5</sup>Kelly (2005), p. 180.

<sup>6</sup>Christensen (2007), p. 193, for the example.

<sup>7</sup>Kelly (2005), p. 180.

weighing the evidence incorrectly on this occasion? What could  $e_2$  be? It cannot be that they are not as epistemically virtuous as you, nor that they do not possess all the evidence that you do, because by hypothesis you are justified in believing they are your epistemic peers in the broad sense. Without some  $e_2$ , your judgment that they have misjudged the evidence is not justified. And then how can you rationally maintain your belief that  $P$ , anymore than either of the mathematical peers in Christensen's example can rationally maintain that their answer as to what each person owes is the correct one?

Well, suppose that  $e_2$  contains a philosophical argument that  $S_1$  is the correct standard, or scale, for weighing evidence, and you have empirical evidence that you have satisfied  $S_1$  and the others have not. Are the others aware of that argument and the relevant empirical evidence? If not, they are not your epistemic peers in the broad sense we are employing, and so that case is irrelevant to the question we are considering. Suppose they are aware of the philosophical argument for  $S_1$ , and the relevant empirical evidence that you have applied  $S_1$  correctly in this case, and yet disagree about whether they have weighed the evidence,  $e_1$ , for  $P$  incorrectly. This may be either because they do not believe that  $e_2$  supports  $S_1$  or because they do not believe that the empirical evidence supports your view that you have applied  $S_1$  correctly. How can you rationally maintain your belief that  $P$  in light of this disagreement by your epistemic peers and rationally maintain your level of confidence in  $Q$  = they have misjudged the evidence,  $e_1$ , for  $P$ ? Assume that they have heard your arguments for  $Q$  yet remain unpersuaded. Why think that you are right in thinking that  $S_1$  is the correct standard for weighing evidence like  $e_1$  and that it is not rather their standard,  $T_1$ ?

Well, let's try one more iteration. Suppose you have another philosophical argument that constitutes evidence  $e_2^*$  for some standard,  $S_2$ , and you then use  $S_2$  to argue that  $S_1$  is the correct

standard for weighing the evidence,  $e_1$ , for  $P$ . Suppose, also, that you have empirical evidence that you have applied  $S_2$  correctly in concluding that  $S_1$  is the correct standard for weighing  $e_1$ . Now either the others are aware of  $e_2^*$ , and the relevant empirical evidence about your application of  $S_2$ , or they are not. If they are not, they are not your epistemic peers in the technical sense being used here, and we can ignore them. If they are, assume that they still disagree about whether  $P$  is true or not. They think that  $e_2^*$  supports  $T_2$ , and  $T_2$  supports  $T_1$ , and using  $T_1$  to weigh  $e_1$  yields the conclusion that it is rational to believe some not- $P$ . So why shouldn't your confidence in  $P$  lessen? Your epistemic peers are unconvinced by your philosophical arguments for various principles to weigh the evidence. It does not matter whether you in fact subscribe to the correct principles ( $S_1$ ,  $S_2$ , etc.). What matters from the standpoint of rationality is *whether you have sufficient reason to think* you subscribe to the correct principles, and the fact that you have not been able to convince your epistemic peers is some reason to think your principles are *not* the correct ones. You better lower your confidence in your principles even if it's true that *if* they are the correct principles, *then* the evidence supports  $P$  more strongly than some non- $P$ . And if you better lower your confidence in your principles, you better lower your confidence in what the evidence supports *given those principles*. Contra Kelly, rational belief is not simply a function of what the correct weighing principles support on the evidence.<sup>8</sup> It's a function of what you have reason to believe are the correct weighing principles, plus your reasons to believe that on them the evidence supports  $P$ . The fact that your epistemic peers disagree about what the evidence supports gives you some, though not necessarily conclusive, reason to believe that you do not possess the correct weighing principles, or that you have applied the correct principles incorrectly on the particular occasion.

---

<sup>8</sup>Feldman (2006) makes a similar point, pp. 230-33, the section on "Objective Evidential Support."

So far I have been discussing the question of whether known disagreement by known epistemic peers rationally requires you to lower your confidence in what you believe on the evidence. I have argued that it does. I now want to turn to a different question. Suppose, taking all the evidence into account, including disagreement by your epistemic peers, it is rationally permissible for you to believe *P*. Can it also be rationally permissible for you to believe some non-*P* on that evidence? I'm not asking whether it can be rationally permissible for you to believe *both* *P* and not-*P* on the evidence. I assume that that it is not rationally permissible, at least if you realize the *P* and some given non-*P* are contradictory. What I am asking is whether it is possible for it to be rationally permissible for you to believe some non-*P* *instead of* *P* on the evidence, given that it is rationally permissible for you to believe *P* on that evidence.

Peter vanInwagen has argued that it cannot be rationally permissible for you to believe *either* *P* *or* some non-*P* on the same evidence. Why not? He offers the following principle:

[V1] If it is rational for a person to accept a certain proposition, it cannot also be rational for that person (at the same time, in the same circumstances) to accept its denial [instead].

His defense of (V1) is the following, where *E* represents the person's total evidence:

It *cannot* be that it would be rational to accept *p* on the basis of *E* *and* rational to accept the denial of *p* on the basis of *E*. Loosely speaking, a body of evidence cannot have the power to confer rationality on both a proposition and its denial. And, finally, it is rational for one to accept a proposition at a certain moment if and only if one's total evidence at that moment bears the impersonal "confers rationality on" relation to that proposition.<sup>9</sup>

In this argument the first occurrence of "rational" in the first sentence either means

---

<sup>9</sup>van Inwagen, "We're Right, They're Wrong," p. 8 in MS.

“rationally required” or “rationally permissible,” and the second occurrence means “rationally permissible.” There is no question that a person who is rationally *required* to believe P on the evidence is rationally *prohibited* from believing some non-P instead. So vanInwagen must think that being rationally *permitted* to believe P on E implies that you are rationally *prohibited* from believing some non-P on E instead. But that is just what (V1) asserts. vanInwagen’s “argument” just reasserts (V1) and does not explain *why* it is true. Why can’t it be rationally permissible for you to believe either P, or some non-P instead, on E? Why can’t a body of evidence “confer rationality on both a proposition and its denial”?

Think of a moral analogy. Suppose I have promised Jones always to tell him the truth, but that he is only an acquaintance, not a good friend. On the other hand, Smith is a good friend. Suppose Jones asks me something about Smith, say, whether he steals petty cash that has been left out or whether he is sleeping with Jones’ girlfriend. If I tell Jones what he wants to know, I betray Smith. If I don’t tell Jones, I break a promise to him and, let’s assume, he will be harmed in some minor way by not knowing the truth about Smith (say, he’ll lose a small amount of money that Smith will steal when he visits Jones or he’ll contract a cold that Jones’ girlfriend has got from Smith). Some people may think that betrayal weighs more heavily on the moral scales and conclude that I should break my promise to Jones and not tell him the truth. Others will think that the promise and preventing minor harm to Jones weigh more heavily and conclude that I should tell Jones the truth about Smith. It’s plausible to think that both weighings of the moral considerations are legitimate, though one implies that I ought *not* tell Jones the truth and the other that I should. In circumstances where morality does not require that I weigh the relevant considerations one way rather than another, it is morally permissible for someone to act in either of the ways, either doing A or not doing A.

Why can't something similar hold when it comes to weighing the evidence for and against some proposition? Suppose we are in a situation where what we should believe is the hypothesis that best explains the observations. Suppose one hypothesis is simpler than another, but that the other hypothesis is a better fit with the data. What is the best explanation of the observations is a function of both simplicity and fit with the data. Why couldn't it be rationally permissible to either weigh simplicity, or fit with the data, more heavily? <sup>10</sup> If so, then it could be rationally permissible for me to believe either of two competing hypotheses.

Richard Feldman disagrees.<sup>11</sup> In a case that I will adapt slightly involving an unknown fork in the road, he says that when I come upon the fork, it can be reasonable for me *to take* either path if I have no evidence that one will more likely lead me to my destination than the other. Despite that, he thinks that it cannot be reasonable *to believe* that fork A will lead me to my destination or reasonable to believe instead that fork B will. He writes,

Believing differs from acting, in a case like this. The reasonable attitude to take toward the proposition that, say, the left path is the correct path is suspension of judgment. Neither belief nor disbelief is supported.... As this case illustrates, acting and believing are different. Sometimes it is reasonable to act a certain way while it is not reasonable to believe that that way of acting will be successful.<sup>12</sup>

Feldman considers another two-person case where the total evidence points as much to Lefty as to Righty being the perpetrator of some crime. He thinks it cannot be reasonable for one person to believe that Lefty, and another that Righty, did it on the same evidence. This leads him

---

<sup>10</sup>David Christensen (2007) considers a case like this, pp. 190-91, and concludes that if I know that an epistemic peer weighs the factors differently I should either believe that she weighs them incorrectly or at least "move my belief in the direction of hers."

<sup>11</sup>And let's face it, he is my epistemic superior!

<sup>12</sup>Richard Feldman, "Reasonable Religious Disagreements," in *Philosophers Without God: Meditations on Atheism and the Secular Life* (Oxford: Oxford University Press, 2007), Louise M. Antony, ed., pp. 194-214; case and quote

to formulate what he calls “The Uniqueness Thesis”:

[UT] This is the idea that a body of evidence justifies at most one proposition out of a competing set of propositions (e.g., one theory out of a bunch of exclusive alternatives), and that it justifies at most one attitude toward any particular proposition. As I think of things, our options with respect to any proposition are believing, disbelieving, and suspending judgment. The Uniqueness Thesis says that, given a body of evidence, one of these attitudes is the rationally justified one.<sup>13</sup>

vanInwagen accepts a principle similar to the Uniqueness Thesis, a principle that is implied by that Thesis. His principle is:

[V2] If two people have the same evidence, and if one of them accepts a certain proposition and the other accepts its denial, at least one of them is not rational: either it is not rational for the one to accept the proposition, or (inclusive) it is not rational for the other to accept its denial.<sup>14</sup>

Perhaps it can *never* be rationally permissible for a *single person* to believe either P or some non-P instead because a given scale to weigh evidence cannot tip both in the direction of some proposition and away from it at the same time. However, why isn’t it possible for there to be at least two different scales that weigh evidence differently, but where there is no fact of the matter about which scale is the correct one? Maybe in the case involving the fork in the road, and the one involving Lefty and Righty, it is not possible for there to be two different scales of this sort. But why not when it comes to weighing factors that are relevant to determining whether one hypothesis better explains some phenomena than another? So I might be rational in believing H1 on total

---

from p. 203. See, also, Feldman (2006), p. 229 where he considers the same case.

<sup>13</sup>Feldman (2007), p. 205. See, also, Feldman (2006), pp. 228-29 and p. 234 for cases where clearly rationality requires suspension of judgment, just as in the Lefty/Righty case.

evidence E and you in believing H2 on that same evidence, where H1 and H2 cannot both be true. The analogy involving the weighing of conflicting *prima facie* duties makes it seem possible that people could rationally weigh differently the factors that determine whether some hypothesis is better than another. In the moral case, it followed that in the relevant situation it was permissible either to do A or not do A. So why not something similar in the epistemic case? If so, there will be cases where it is epistemically permissible for one person to believe P and the other to believe some non-P, even though they have the same total evidence.

Further, there seem to be examples where it is rational for me to believe one thing on the evidence and rational of you to believe the opposite. Suppose you think that the Detroit Tigers will win their division, and I disagree. Your reasons might be that the pitching is so good that it will carry them through even though Magglio Ordonez and Placido Polanco are off in their hitting. I might counter that the pitching isn't that good. They still need left-handed pitching and Dontrelle Willis is very shaky, despite a good first game on the comeback trail, and Bobby Seay's pitching out of the bullpen has been less than stellar. After a good start, Armando Galarraga's pitching has been atrocious. We might have opposite views about the Tigers' chances of winning their division even though we share the same evidence. We know everyone's batting and on-base averages, the won-loss records and the earned-run averages (ERAs) of all the pitchers, whom the Tigers might call up from the minors, etc. Still, we disagree. The best explanation seems to be that rational people can weigh the relevant considerations differently, including, in this case, the quality of the pitchers and their performances, and the importance of pitching as compared to hitting in winning ballgames.

Consider another similar case. Suppose we are picking which horse will win the Belmont

---

<sup>14</sup>vanInwagen, "We're Right. They're Wrong," MS, p. 9.

Stakes. We both know the racing histories of the horses entered, their track records and bloodlines, and the successes and failures of the jockeys who will ride them. I pick Mine That Bird who will be ridden by Calvin Borel, who rode the winners in both the Kentucky Derby and the Preakness; you pick the big gray horse, Dunkirk, who will be ridden by John Velazquez. Assume that we both have the same evidence and similar track records at picking winners and we both know this. Isn't it rationally permissible for me to believe that Mine That Bird will win and for you to believe it will be Dunkirk?

Feldman argues that the idea that either of two ways of weighing factors can be reasonable "...just pushes the question back a step. We can now ask which factor *should* be weighed more heavily."<sup>15</sup> Feldman thinks that people need reasons for weighing evidence the way they do, that they do not get "their preferred way of weighing the factors" for "free."<sup>16</sup> But if at some point they do not, an infinite regress looms. As we saw earlier, maybe I can justify accepting scale, S1, by assuming scale, S2, in conjunction with evidence, E1. And given S1, I can justify believing P on evidence, e1. But if I have to have reasons to accept S2, won't I have to assume some scale S3 that, in conjunction with, say, E2 justifies S2? Where will this stop? It looks as though to be justified at all at some point I must rely on a scale that I have no reason to accept, but that, nevertheless, is one that provides justification when the appropriate evidence is placed on it.

There is reason to think that certain scales do not produce reasonable belief when evidence

---

<sup>15</sup>Feldman (2007), p. 206; my italics.

<sup>16</sup>Feldman (2007), p. 206. In (2006), p. 232, Feldman argues that there is not a *general* requirement for justification that you have evidence that your evidence is good, but there is once you know that known epistemic peers disagree, for that constitutes and undermining defeater. Feldman (2006), p. 226, writes of a person who has different starting points from another, "...if both alternatives are reasonable, and he lacks a reason to prefer his own alternative, then accepting that alternative rather than the other is arbitrary. It is difficult to see how anything 'downstream' from this arbitrarily selected alternative counts as justified. It is therefore implausible to think that he can regard the disagreement as a mutually reasonable one." But why not? Couldn't it be reasonable for you to think that a particular color of car is the best color and reasonable for someone else to think it is some other color (as long as no one thinks it's pink and yellow, say!)?

is placed on them, say, a counterinductive scale that implies that if all the crows you've seen are black, then it's reasonable to believe that the next crow you see will *not* be black, or a view about inference to the best explanation that implies that the best explanation of the totality of our current evidence is that the earth is flat.<sup>17</sup> Still, it is reasonable to believe that some scales for weighing evidence *do* yield reasonable beliefs, and there is no reason to think that *only one* such scale can produce reasonable beliefs. In fact, the examples involving baseball and horse racing suggest that more than one scale can count as a reasonable way of weighing the factors.

So is Tom Kelly right after all? Does the rationality of a belief depend solely on how the evidence for that belief weighs on the correct epistemic scale? Earlier I objected that just judging that your epistemic peer has misjudged the probative force of the evidence is not enough for you to legitimately discount known disagreement with him. I said you would need some reason that you have and he lacks to legitimately do that, and then, of course, he would not be your epistemic peer in the broad sense. But there I was considering whether it could be rational of you *not to lower your confidence in what you believe, for no reason*, on known disagreement with a known epistemic peer. I argued that it could not be. Here I am arguing that *it could be rational* for you to believe P, and rational for your epistemic peer to believe some non-P, given the same total evidence. That will be the case if, and only if, you in fact weigh the evidence on scales that are equally good from an epistemic standpoint. But the fact that you weigh the evidence differently from your known epistemic peer is not enough for it to be rational for you *not to lower your confidence in what you believe* and your peer denies.

Up to this point I have argued that known disagreement with a known epistemic peer

---

<sup>17</sup>vanInwagen, "Freedom to Break the Laws," MS, note 13 thinks that a person whose belief that the earth is flat is in equilibrium with his other beliefs is not his epistemic peer, presumably because he lacks some epistemic virtues, including ones that would prevent him from weighing the evidence in a way that comes out supporting the flat earth

rationally requires you to lower your confidence in what you believe, but because the Uniqueness Thesis is false, it permits both of you to be rational in what you believe. I now want to turn to some philosophical disputes and see whether the correct diagnosis is that the disputants are epistemic peers that weigh the evidence differently because they use different, but equally good, epistemic scales.

In several essays, Peter vanInwagen has noted the disagreement between him and David Lewis over whether incompatibilism or compatibilism is true regarding the relation between determinism and free will. vanInwagen is an incompatibilist and holds (I) if determinism is true, no one ever has free will. Lewis was a compatibilist and denied (I); he thought it possible for determinism to be true *and* for people to have free will. Recall that vanInwagen accepts the following:

[V2] If two people have the same evidence, and if one of them accepts a certain proposition and the other accepts its denial, at least one of them is not rational: either it is not rational for the one to accept the proposition, or (inclusive) it is not rational for the other to accept its denial.

vanInwagen thinks the (V2) forces him into accepting some unacceptable conclusion. Either he has to accept (a) that he and Lewis did not have the same evidence, (b) that they did not have the same epistemic virtues, (c) that Lewis was not rational in accepting compatibilism, or (d) that he, vanInwagen, was not rational in denying compatibilism and accepting incompatibilism. He thinks that if evidence in philosophy just consists of arguments, distinctions, analyses, and examples, then it is not plausible to think that they did not have the same evidence. On the other hand, it is not plausible to think that either he or Lewis was being irrational, nor even less than fully rational if

---

hypothesis. Feldman (2007), p. 211 and (2006), p. 230 makes a similar point about those who believe in astrology and

being irrational requires you to be *far less* than fully rational.

vanInwagen implicitly recognizes that another clause about epistemic virtue must be added to the antecedent of (V2) for it to have even a chance of being true, for one might hold that if two people have the same total evidence it does not follow that at least one must not have a rational belief since they may not have equal epistemic virtues despite having equal evidence. The idea is that a person with diminished epistemic virtue might rationally believe some non-P on the total evidence, while anyone with more epistemic virtue would believe P on that same evidence, but that the lower epistemic virtue gives the first person a legitimate excuse so that he is rational in believing what he does, that is, is epistemically blameless in believing what he does. But as vanInwagen says, “One could hardly hold that David was stupid or lacking in philosophical ability or that he labored under any other cognitive deficiency relevant to thinking about the problem of free will.”<sup>18</sup> So even an amended (V2) seems to leave us with no escape from the unacceptable alternatives.

Finding all the options that he is faced with unacceptable if he accepts (V2), vanInwagen sees that some of the principles that led him to (V2) must be wrong. But he finds it hard to give up those principles or (V2) itself. So he toys with the idea that not all evidence is public, that some cannot be expressed in language, is interior and incommunicable, and cannot be passed from one person to another. After giving an argument that says that evidence cannot both direct a person toward some belief and away from it, vanInwagen says, “The difficulty in finding anything to say in response to this argument, taken together with my unwillingness to concede either that I am irrational in being an incompatibilist or that David was irrational in being a compatibilist, tempts

---

credits vanInwagen (1996) with making the same point earlier.

<sup>18</sup>vanInwagen, “We’re Right, They’re Wrong,” p. 11 in MS. See pp. 10-11 for the basis of my paragraph in the text. He makes similar points in “Listening to Clifford’s Ghost,” pp. 10-15 in a 25 pp. MS.

me to suppose that I have some sort of interior, incommunicable evidence (evidence David did not have) that supports incompatibilism.”<sup>19</sup> vanInwagen confesses that although this view is tempting, “[it] is hard to believe” because it is hard to believe that in the many cases in which he disagrees with other philosophers, he always possesses “interior, incommunicable” evidence that they lack.

20

He thinks it more plausible to believe that *both* he and Lewis were rational in holding opposite views *on the same evidence*, but he cannot see how that could be. He assumes that for that to be possible the evidence would have to direct him towards incompatibilism and Lewis away from it at the same time. vanInwagen also assumes that evidence “tracks the truth.” But how could it do that if it could direct people both toward and away from some proposition? How could it make it more likely that some proposition is true and more likely that its denial is true? vanInwagen is left in a quandary.

One thing that puts him in this quandary is the assumption that evidence tracks the truth in the sense of making it objectively more likely that the proposition that it is evidence for is true. It does not necessarily do that as demon-world and Matrix examples show. Of course, it is true that if you have good evidence for P then there is *reason to believe that P is true*, but that is compatible with its being objectively unlikely that P is true. So that makes it possible for evidence to point both ways. But even then, it does seem odd to say that evidence E can make it reasonable to believe P and reasonable to believe some non-P.

---

<sup>19</sup>vanInwagen, “We’re Right, They’re Wrong,” p. 13 in MS. See, also, “Listening to Clifford’s Ghost,” pp. 18-19.

<sup>20</sup>Feldman (2006) takes up a similar suggestion in his section “Private Evidence,” pp. 222-24, and concludes that insofar as epistemic peers know about each other’s “insight,” each has evidence that the other has (private) evidence. So why think that your private evidence is good but the other person’s is not? I think that these considerations ultimately undermine Michael Bermann’s attempt (forthcoming in *Episteme*) to show that people of equal epistemic virtue, and who share all the evidence it is possible to share, can, nevertheless, be rational in believing opposite things based on different insights and different error theories that they use to explain away the evidential force of the “insights” of their opponents.

My diagnosis of where vanInwagen has gone wrong is in adopting the metaphor of evidence *directing one toward some proposition*. That metaphor makes it hard to see how evidence could count for some proposition and its denial. My metaphor of evidence *weighing for or against some proposition* seems more apt. It enables us to explain how two people with the same total evidence can be fully rational in believing opposite things since they will be fully rational if they weigh the evidence on different scales and the scales are equally good from the standpoint of epistemic rationality. My solution enables us to see how it is possible for one excellent philosopher to rationally accept incompatibilism, and another to rationally accept compatibilism, *on the same evidence* (where interior and incommunicable “evidence” doesn’t count as evidence) *provided* they weigh the evidence on scales that are equally good from an epistemic standpoint.

Do I think that this is usually what happens when excellent philosophers disagree? No! I think that almost always they have different evidence. That is because philosophers are good at thinking of objections to what they have just read or heard, those thoughts constitute evidence that bears on the issue at hand, and often the opposition is unaware of those thoughts. This is interior, though not incommunicable, evidence. If Feldman’s terminology (2006), they are often justified in isolation and the situation is not one of full disclosure.

Here’s an example. I am not an act utilitarian because I think there are examples of using people as mere means that maximize utility but that, nevertheless, are wrong. However, I am bothered by Peter Unger’s multiple-option trolley case which concludes that an act that I would characterize as using someone as a mere means is not wrong. Here are the options: (1) let a runaway trolley run over six people who cannot get off the track the trolley is on; (2) turn the trolley down a spur track where only three are trapped; (3) run an empty side trolley into the

runaway trolley that now has two passengers on it, knocking it off the track before it reaches the six but killing the two passengers; (4) same as (3) but now the passengers are on the side trolley and they will be killed by the collision, but now no one is on the runaway trolley; (5) run a heavy man on roller skates into the empty runaway trolley killing him but knocking the trolley off the tracks, thereby saving the six. Unger argues that (2) is morally permissible and if (2) is, (3) is; if (3) is, (4) is, and if (4) is, (5) is. So (5) is morally permissible, even morally obligatory, since fewer innocent people are killed in (5) than in any of (2)-(4)—and this despite the fact that the heavy man is being used as a mere means to save the six.<sup>21</sup>

Last semester one of my students<sup>22</sup> pointed out that only in (5) is it true that someone's body is essential to saving the others, and in that regard, (5) resembles Transplant, where to save five you must cut up one and distribute his vital organs to the five. The idea is that using someone's body without his consent is very wrong, and constitutes using that person as a mere means (even if other things can also qualify as using someone as a mere means). That difference can explain the difference between (5) and all the other options (2)-(4). Of course, if the weight of the passengers were essential in (4) to knocking the runaway trolley off the tracks, then perhaps (4) would not be morally different from (5) on the suggestion being considered. But maybe people would have the intuition in new-(4) that it is wrong to run the side-trolley into the runaway trolley. In any case, my point is that perhaps Unger did not think of this difference between (4) and (5) when he was constructing his example. If he did not, then he does not have the same evidence that my student, and now I, have.

What about the disagreement between vanInwagen and Lewis? Didn't they have the same

---

<sup>21</sup>Peter Unger (1996) gives the this multiple-option case in Chpt. 4, pp. 84-118. A drawing of the options by Jesse Prinz is on p. 89, though it shows option (4), not (3).

<sup>22</sup>Tiffany Hudson.

evidence? And isn't it implausible to think that they were using different epistemic scales to weigh that evidence? I'm not sure that they would really disagree if they considered the following argument which employs what van Inwagen calls the "Consequence Argument" as a sub-argument for the first premise.

1. If determinism is true, then no wholly human creature has free will (for if they did, they would have to be able to falsify a law of nature, that is, to perform a miracle, and no wholly human person can do that. Maybe Jesus could, but he was not wholly human. This is the Consequence Argument.).
2. If no wholly human creature has free will, then no such creature is ever morally responsible for anything he does.
3. Therefore, if determinism is true, then no wholly human creature is morally responsible for anything he does.

Lewis will deny that (2) is true when "free will" is understood in the way that vanInwagen thinks makes (1) true. He will deny that if no human creature can perform a miracle, then no such creature is ever morally responsible. Instead, he will say that there is a sense of "free will" (having to do with being coerced, drugged, hypnotized, brain washed, etc.) according to which (2) is true. In that sense of "free will," it is not true that free will requires the ability to perform a miracle. Still, that sort of free will is enough for moral responsibility. So Lewis will say that the argument equivocates on "free will." The sense that makes (2) true makes (1) false, and vice versa. Perhaps vanInwagen is right in thinking that for there to be free will if determinism is true, people must be able to perform a miracle. But that sort of free will is not required for moral responsibility. So the sense of "free will" that makes (1) true makes (2) false. No doubt vanInwagen will disagree, saying instead that the sense that makes (1) true also makes (2) true. He will say that people need

free will in the strong “ability-to-perform-a-miracle” sense for there to be moral responsibility, not the weak, little wimpy sense of being uncoerced, etc., that the compatibilists propose.

Now I think the correct response to vanInwagen is to say that there is also an equivocation on “is morally responsible” (as well as on “free will”). It could mean “is morally responsible from the standpoint of justifiably imposing sanctions” or it could mean “is really worthy of blame.” I would think that there is a compatibilist sense of “free will” that implies that a person who has it is morally responsible in the first sense, that is, in the sense that she can justifiably *be held responsible* from a moral standpoint, but that denial of free will in the compatibilist sense is not implied by determinism. So the compatibilist can argue that it is possible for determinism to be true and for there to be free will in a sense relevant to moral responsibility. Hence, it is possible for determinism to be true and for people to be justifiably *held* morally responsible for some of their actions. However, the compatibilist should grant that *in vanInwagen’s sense of free will*, determinism implies the lack of free will, *and* in that sense, *no one is really morally blameworthy* for anything they do. If the only way that Bill could have failed to kill Jill was by performing a miracle, then Bill is not really blameworthy for killing Jill. Still, the good consequences of *holding* Bill responsible, can justify imposing sanctions, even if, strictly speaking, it cannot justify imposing punishments on him because it is conceptually necessary that punishment is justifiable only if the person is really morally blameworthy.

Only this “double equivocation” reading of the argument for incompatibilism promises to reconcile the incompatibilist and the compatibilist. Before noticing this double equivocation, both could have been rational in accepting what they did in the sense that neither was *grossly irrational* in believing what they did, or that neither was epistemically blameworthy for believing what they did, even though neither was perfectly rational since their views rested on an equivocation. Once

that equivocation has been revealed, then I think the rational thing for both to believe is: (a) if determinism is true, then no one is really worthy of blame (a la the incompatibilist) and (b) even if determinism is true, some can be justifiably sanctioned for their actions, the justification stemming from the good consequences, in terms of deterrence and reform, that will result if certain people *are held responsible* for their actions.

If less than perfect rationality is possible when subtle equivocations are involved, then there is another way to consistently hold that two epistemically virtuous people with opposite views can rationally hold them on the same evidence, while still acknowledging that in some sense that seems impossible. The sense in which it seems impossible is if being rational is understood to mean being *fully or perfectly* rational, while it seems possible if being rational is understood to mean not falling too short of full or perfect rationality, that is, falling somewhere in the range between perfect rationality and irrationality.<sup>23</sup> My solution says that it is possible for epistemic peers who disagree to be fully or perfectly rational on the same evidence because it is perfectly rational for them to use the scales they do and the different scales weigh evidence differently. The alternative solution implies that that is not possible, but grants that some sort of less than full or perfect rationality is possible. Either view frees vanInwagen from his quandary. I hold my view on the basis of the analogy with ethics where in some cases there is not just one correct way to weigh the relevant *prima facie* duties, or prohibitions.

So far I have considered disagreements in moral philosophy and on the free will/determinism question. I want to consider disagreements in one more area, namely, in philosophy of religion. There are many excellent philosophers who are theists and many who are

---

<sup>23</sup>My colleague Larry Powers has a developed metaphilosophy in which equivocation plays a central role, and he allows for two different senses of “rationality.” So I had his metaphilosophy in mind when I thought of the alternative solution. I am also indebted to a conversation with him that made me realize that there may be a “double”

atheists. There are so many arguments, objections, replies, counter-replies, etc., in this broad field that it is doubtful that any two philosophers have exactly the same evidence regarding the question of God's existence. So the question of whether two epistemic peers in the broad sense, that is, those who have the same epistemic virtues *and* the same evidence, can rationally believe opposite things about God's existence, does not arise. But I am bothered by narrower claims in this area where the question does arise. For instance, I do not see how theists can rationally believe in a God who is all-knowing, all-powerful, and wholly good in light of the problem of evil. But I want to focus on even narrower claims.

Alvin Plantinga thinks that propositions such as: (P1) God created this flower; (P2) God created this sunset; (P3) God disapproves of what I have done, and so on, can be properly basic for some people. By properly basic he means that a person can be warranted in believing them not on the basis of any evidence. Given what he means by "warrant," which he takes to be that which must be added to true belief to turn it into knowledge, it follows that a person could know these things if they are true. Since (P1) and (P2) are causal claims, they are like the claim (Q1) vanGogh painted this painting or (Q2) Frank Lloyd Wright designed this building. Now I think everyone would agree that (Q1) and (Q2) are not properly basic for anyone. Testimonial evidence aside, a person would have to have evidence stemming from his observations of other known vanGoghs, or structures known to be designed by Wright, to have warrant for believing that a particular painting or building was by vanGogh or Wright, respectively. And testimony would provide evidence only if there was further evidence to think the testimony is reliable. While beliefs in the relevant propositions could be *psychologically* basic for experts because they would form them immediately on viewing the vanGogh or the building designed by Wright without having to infer

them from anything else, they could not be *properly* basic for anyone, not even the experts.

Everyone needs some background evidence to be warranted, or justified, in believing the relevant propositions.

If that is true of (Q1) and (Q2), it is reasonable to believe that it is also true of (P1) and (P2). Similar arguments can be constructed concerning (P3). So Plantinga should not believe that any of (P1)-(P3) are properly basic for anyone. But he thinks all of them can be for certain persons! No one can deny that Plantinga is an excellent philosopher. Is it that he never thought of the parallels between the Qs and Ps? Possibly, and this is the most charitable and friendly interpretation. The other one is that on this particular occasion his epistemic virtues were not fully engaged and so his belief that (P1) and (P2) can be properly basic for some people is not rational, just as occasionally a morally virtuous person will do the wrong thing.

Plantinga also gives an example that seems to me to be a clear counterexample to his general account of knowledge. The example involves a strange orange bird that lives on another planet whose call sounds like church bells. We are to assume that the inhabitants of that planet have the concept of orange and of sounding like church bells ringing but have no evidence that sounding like church bells ringing is a reliable sign that something orange is nearby. Nevertheless, Plantinga assumes that the inhabitants are designed (by God or evolution) in such a way that when they hear this sound they immediately form the belief that there is something orange nearby, and there is. On this planet, sounding like church bells ringing *is* a reliable sign that something orange is nearby. Plantinga asks rhetorically of the planet's inhabitants, "Why wouldn't this belief have a high degree of warrant—indeed, why couldn't it constitute knowledge?"<sup>24</sup> Of course, his view

---

dissertation, *Knowledge and Meaning in Philosophy*, Cornell University Ph.D dissertation, January, 1977.

<sup>24</sup>This quote and example are taken from Alvin Plantinga, *Warrant: The Current Debate* (Oxford: Oxford University Press, 1993), pp. 62-63.

implies that the belief has warrant and constitutes knowledge since their belief is founded on a cognitive mechanism functioning properly in an appropriate environment according to a design plan successfully aimed at truth. In the setting of their planet, the design plan of a cognitive mechanism in the planet's inhabitants yields true beliefs about color upon hearing the church-bell sound.<sup>25</sup>

This example is relevantly like other examples used against reliabilists about justification, or reliabilists who hold that knowledge is reliably produced true belief, in that in all the cases the believers have no evidence that what they believe is true or that they were reliably produced. How can Plantinga think that this example is not a counterexample to his theory of warrant and knowledge; how can the reliabilists think that Lehrer's Truetemp and Bonjour's Norman, the clairvoyant, are not counterexamples to their theories? I think they are irrational in denying they are, just as they would be if they denied that certain Gettier examples (say, Chisholm's example involving the fake sheep in the field) are counterexamples to the justified true belief account of knowledge.

---

<sup>25</sup>I discuss Plantinga's example in "Moral and Epistemic Duty," in *Knowledge, Truth, and Duty*, Matthias Steup (ed.), (Oxford: Oxford University Press, 2001), pp. 34-48; see, pp. 45-56 for the discussion of Plantinga and this example.

In most philosophical disagreements among philosophers who possess the same epistemic virtues, the disputants either do not have the same evidence or at least one of them is not being fully rational on the occasion despite her epistemic virtues. While it is possible that they both have the same epistemic virtues and the same evidence (that is, are epistemic peers in the broad sense), and that both are fully rational on the particular occasion, it is rare. So even if, unlike me, you hold that if two people are epistemic peers in the broad sense, then it is not possible for them both to be fully rational in believing opposite things, you will not be forced to conclude that, in philosophy, it is often true that at least one of the epistemic peers is not fully rational, for often they will not have the same evidence, despite appearances to contrary. Of course, I go further and hold that often in philosophical disputes at least one of the epistemic peers in the broad sense is not being fully rational on the particular occasion. Often at least one, and probably both, are misled by subtle equivocations, and sometimes at least one of the epistemic peers is irrational in his assessment of the evidence on the particular occasion.<sup>26</sup>

Bruce Russell  
Department of Philosophy  
Wayne State University  
Detroit, MI 48202

[bruce.russell@wayne.edu](mailto:bruce.russell@wayne.edu)

---

<sup>26</sup>I sometimes call myself an unfriendly atheist since I hold that theists often fail to be fully or perfectly rational in disputes about God's existence. In "Listening to Clifford's Ghost," p. 24, vanInwagen thinks that someone is "a comic figure" if he thinks that he has sufficient evidence for his beliefs on some substantive philosophical question but that those philosophers who disagree with him lack some epistemic virtue or exhibit some cognitive or epistemic defect! I do not go so far as to say they lack some epistemic virtue, but I do say that they exhibit some epistemic or cognitive or epistemic defect on the particular occasion, or have not thought of some relevant considerations.

## References

- Audi, Robert ( ). “Intuition, Inference, and Rational Disagreement in Ethics.” *Ethical Theory and Moral Practice*: 18 pp. MS.
- Bergmann, Michael (forthcoming). “Rational Disagreement After Full Disclosure.” *Episteme*, 24pp. in MS.
- Christensen, David (2007). “Epistemology of Disagreement: The Good News.” *Philosophical Review*, 116(2): 87-117.
- Feldman, Richard (2006). “Epistemological Puzzles About Disagreement,” in Stephen Hetherington (ed.), *Epistemology Futures*. Oxford: Oxford University Press, 216-36.
- \_\_\_\_\_ (2007). “Reasonable Religious Disagreements,” in Louise M. Antony (ed.), *Philosophy Without God: Meditations on Atheism and the Secular Life*. Oxford: Oxford University Press, 194-214.
- Kelly, Thomas (2005). “The Epistemic Significance of Disagreement,” in Tamar Szabo Gendler and John Hawthorne (eds.), *Oxford Studies in Epistemology*. Oxford: Oxford University Press, 167-96.
- Plantinga, Alvin. (1993). *Warrant: The Current Debate*. Oxford: Oxford University Press.
- Powers, Lawrence (1977). *Knowledge and Meaning in Philosophy*. Ph.D dissertation, Cornell University.
- Russell, Bruce (2001). “Epistemic and Moral Duty,” in Matthias Steup, ed., *Knowledge, Truth, and Duty*. Oxford: Oxford University Press.

vanInwagen, Peter (1996). "It Is Wrong Everywhere, Always, and for Anyone To Believe Anything on Insufficient Evidence," in Jeff Jordan and Daniel Howard-Snyder (eds.), *Faith, Freedom, and Rationality: Philosophy of Religion Today*. Lanham, MD: Rowman & Littlefield, 137-53.

\_\_\_\_\_ (2004). "Freedom to Break the Laws." *Midwest Studies in Philosophy* 28: 334-50.

\_\_\_\_\_ (forthcoming, a). "Listening to Clifford's Ghost." 25 pp. in MS.

\_\_\_\_\_ (forthcoming, b). "We're Right. They're Wrong." 15pp. in MS.

Unger, Peter (1996). *Living High & Letting Die*. Oxford: Oxford University Press.