

## Two Notions of Epistemic Rationality

*Abstract:* In this paper I argue that we need two notions of epistemic rationality. We need one notion to help us *deliberate* and another to help us *evaluate*. I motivate this distinction by (a) arguing that there are competing desiderata for a theory of rationality that no single notion can satisfy and (b) showing that the distinction between these two notions can resolve a tension between plausible views about higher order evidence and views about how we should take into account the opinions of others.

### 1. Introduction

We demand too much of the notion of epistemic rationality. I hope to convince you that we need (at least) two notions of rationality to do the work that epistemologists usually try to do with one. I think that there are a number of independent ways to motivate such a distinction<sup>1</sup> but, in this paper, I intend to motivate the distinction by (a) arguing that there are competing desiderata for a theory of rationality that no single notion can satisfy and (b) showing that the distinction between these two notions can resolve a tension within a combination of views that many epistemologists are attracted to. These two strategies are related. Each of the views that are in tension gains its intuitive plausibility from the fact that it does well with respect to one of the desiderata for a theory of rationality. The tension between them arises because no single notion can satisfy *both* of the desiderata. I will begin by briefly describing the views that are in tension, and the distinction that, I think, will save the day.

Here are rough versions of the two views which, I will argue, are in tension:

ROUGH CONCILIATIONISM - What it is rational to believe depends, in part, on what limitations the agent rationally believes she has. In particular, if S rationally believes that she will do no better than chance at judging whether p on the basis of her evidence, S should suspend judgment on p.<sup>2</sup>

ROUGH DEFERENCE - If you know that S has strictly more evidence than you do (that is, S has all the evidence you have and more) and that S's doxastic attitude towards a proposition p is rational, then it is rational for you to adopt S's doxastic attitude towards p.<sup>3</sup>

---

<sup>1</sup> References omitted for blind review.

<sup>2</sup> The word "conciliationism" comes originally from a view in the disagreement debate according to which you should reduce confidence when you learn that a peer disagrees with you (see for example, Elga (2010)).

<sup>3</sup> If you are a permissivist, you can still accept a version of this principle with the added qualification that S must have the same standards of reasoning as you do, or, if you are a Bayesian, the same prior probabilities

Here is a brief summary of the distinction that I will be defending: there are two primary purposes for which we can put the notion of epistemic rationality to work: we can use the principles of epistemic rationality to help us deliberate about whether or not various propositions are true, and we can use the principles of epistemic rationality to evaluate the attitudes of other agents. One of the most important reasons we have to evaluate the rationality of other agents' attitudes is so that we can use such judgments to determine whether or not we should defer to those agents.<sup>4</sup> I am going to argue that the rules that we should use for deliberation (from a first person perspective) are different from the rules we should use for evaluation (from a third person perspective).

In section two, I will describe and motivate the refined versions of CONCILIATIONISM and DEFERENCE. In section three I will argue for the tension between these two claims. In section four, I develop the distinction between two notions of rationality and show how it solves the problem described in section three. I will then tie up some loose ends in section five and conclude with some brief remarks about what underlies the difference between these two notions of epistemic rationality.

## 2. CONCILIATIONISM and DEFERENCE

In this section I will describe and motivate CONCILIATIONISM and DEFERENCE.

### 2.1 CONCILIATIONISM

CONCILIATIONISM is motivated by cases like the following<sup>5</sup>:

DRUG: You have worked through a logic problem whose answer is either  $p$  or  $\sim p$ . You concluded that  $p$  is the answer. You then learn that your evil logic teacher slipped a reason-distorting drug in the coffee you were drinking while doing the logic problem. People who reason through logic problems under the influence of this drug only get the right answer about 50% of the time.

Such cases have been widely discussed and many have found it compelling to say that, in cases like this, you ought to suspend judgment.<sup>6</sup> Conciliationists have explained this judgment by arguing that (a) in cases like DRUG you should think it likely that you reasoned

---

that you have (in other words, for deference to be required,  $S$  must also evaluate evidence the same way that you do).

<sup>4</sup> See Dogramaci (2012) for an illuminating discussion of the importance of this particular role for epistemic rationality.

<sup>5</sup> This case is based on a case from Christensen (2010).

<sup>6</sup> For example, Christensen (2010), Elga (ms.), Horowitz and Sliwa (ms.) and Vavova (ms.). Though for notable exceptions see Lasonen-Aarnio (forthcoming), Kelly (2010), Weatherson (ms.) and White (2010). Also note that, for my purposes, it is fine to think of suspension of judgment as having a 0.5 credence.

incorrectly and (b) you need to assess the likelihood that you reasoned incorrectly in a way that is *independent* of the reasoning in question.<sup>7</sup> This independence constraint is meant to rule out a response to DRUG like the following: “Look! The answer to the logic problem is p and p is what I concluded. Since I got the answer right I must be immune from the effects of the drug.” If CONCILIATIONISM is to yield the verdict that you should reduce confidence in DRUG, it will have to include a clause which rules out judging that you are reliable *on the basis of the reasoning in question*. This sort of clause is called “an independence clause.”

So the refined version of CONCILIATIONISM<sub>R</sub> says:

CONCILIATIONISM: What it is rational to believe depends, in part, on what limitations the agent rationally believes she has. In particular, if, *independently of S’s reasoning about p*, S rationally believes that she will do no better than chance at evaluating some body of evidence, S’s credence in p should equal her prior credence in p.<sup>8</sup>

Now, of course, conciliationists will have things to say about cases in which you will do better than chance, but still not very well. However, this principle, which says nothing about the more complicated cases, is strong enough to run into the troubles I will be describing. So I will set the more complex cases aside.

## 2.2 DEFERENCE

The refined version of ROUGH DEFERENCE has an important qualification. It only tells you to defer to a subject, S, when the additional evidence that S has, and the proposition in question has *no de se content*. Propositions (and evidence, if evidence is propositional) with de se content are propositions about what time it is, where you are, or who you are. Such propositions distinguish not only between ways the world is, but also between the subject’s possible locations in those worlds. I will call evidence and propositions lacking de se content, “de dicto evidence” and “de dicto propositions” respectively. I will explain why these restrictions are necessary in a moment, but first, let me get the correct version of the principle on the table:

DEFERENCE: If you know that S has strictly more de dicto evidence than you do (that is, S has all the de dicto evidence that you have, plus more), S has no less self-locating

---

<sup>7</sup> The independence requirement is described in the context of disagreement in Christensen (2007) and Elga (2007). There is a discussion of this constraint in the context of higher order evidence more broadly in Christensen (2010) and Vavova (ms.).

<sup>8</sup> The judgment that you should suspend judgment in DRUG relies on the assumption that, prior to your calculations, you didn’t have any reason to prefer p to ~p. If you did, then the conciliatory response won’t require you to have less confidence than you initially started with. It will just require that you return to your prior credence in p.

information than you do, and that S's doxastic attitude towards a de dicto proposition p is rational, then you ought to adopt S's doxastic attitude towards p.<sup>9</sup>

Before defending DEFERENCE, let me explain why these qualifications were added. Without the restriction that the evidence be de dicto, the principle would be trivial, and without the restriction that the propositions be de dicto, the principle would be false.

To see why the evidence restriction is necessary note that without that restriction, there would be very few (if any) cases in which an agent S has strictly more evidence than you do, and so very few (if any) cases in which the principle applies. For to have strictly more evidence than you do, S needs to have all of the same evidence that you have, plus more. But S will rarely (if ever) know all of the essentially indexical propositions that you know<sup>10</sup>. For example, suppose I know that I am Mary, that I am mildly hungry, and that I am in the empire state building. Whatever your expertise, I doubt that you know that *you* are Mary, that *you* are mildly hungry and that *you* are in the empire state building.<sup>11</sup> Since, in any normal case, all agents know at least *some* essentially indexical propositions that other agents do not, if the principle is to apply in the cases in which we want it to apply, the principle must kick in whenever S has more *de dicto* evidence than you do.<sup>12</sup>

The second modification to DEFERENCE<sub>r</sub> involved restricting the propositions about which one should to defer to de dicto propositions. While the first modification was necessary to prevent the principle from being trivial, this modification is necessary to prevent the principle from being false. To see why the principle would be false without the restriction to de dicto propositions, consider the following case. Alex has all of the de dicto evidence that Bonny has, plus more. He also believes that he is Alex. Bonny knows that Alex believes that he is Alex and that Alex has more de dicto evidence than she does. Nonetheless, Bonny should clearly not believe that *she* is Alex. What this example

---

<sup>9</sup> As I mentioned in footnote 3, if you are a permissivist you can still accept a version of this principle. Also, note that (b) is added to sidestep worries about Sleeping Beauty (see Elga (2000)) and Arntzenius' (2003) "prisoner in the cell" case. The idea of restricting Reflection in this way is suggested by Horgan (2004) and Elga (2007).

<sup>10</sup> Perhaps John Perry's (1977) case of Lingens and Lauben the amnesiacs is such a case.

<sup>11</sup> You might think that if I know that I am Mary, and you know that you are George, we can still have the same evidence provided that you know that *I* am Mary and I know that *you* are George. If this is your view, then the de dicto restrictions I will be discussing will not be necessary.

<sup>12</sup> You might think that, even with the restriction to de dicto evidence, this principle will rarely have any applications. For very few people even know all of the *de dicto* propositions that I know. For example, surely I should defer to paleontologists about dinosaur related matters even though none of them know the de dicto proposition about the exact location of my birth. These considerations motivate a stronger version of DEFERENCE according to which you should defer whenever an agent has strictly more *relevant* evidence. What is "relevant evidence?" Perhaps evidence which is such that, if the other agent had that evidence, it would be rational for her to revise her attitude. This is an equally important modification for the purposes of capturing the cases that we want DEFERENCE to apply to. However, since my principle is weaker and is all that is necessary for the primary purposes of this paper, I will leave it as is.

illustrates that, in general, it only makes sense to defer to more informed agents about *de dicto* propositions.<sup>13</sup>

I will argue for DEFERENCE by first motivating a special case of DEFERENCE and then showing how the argument for the special case generalizes. The special case of DEFERENCE that I have in mind is one in which S is a future time slice of yourself. The principle, as applied to a future time slice of yourself, I call “Reflection” since it is very similar to the principle so named and defended by Bas van Fraassen (1984).

**Reflection** -If you know that, in the future, you will have strictly more *de dicto* evidence than you do now, you will not lose self-locating information, and that your later doxastic attitude towards a *de dicto* proposition p is rational, then you ought to adopt that attitude now.

Before defending these principles, I should note that there are well known counterexamples to certain versions them. Although the principles I stated are meant to sidestep *some* of the worries about Reflection and DEFERENCE, I would not be surprised if there were others. (Briggs (2009) provides a thorough discussion of counterexamples to Reflection and argues for a version a principle that avoids them). However, for my purposes, all that’s really important is that deferring to more informed rational attitudes is, *in general*, a sensible thing to do, and an important part of our epistemic practice. Thus, when we encounter cases in which it isn’t sensible to defer, some explanation is necessary.

Here are three arguments for Reflection:

(1) *Reflection Follows from Conditionalization* - In his original paper, van Fraassen argued that perfect Bayesian agents, who conditionalize, will satisfy Reflection. It turns out that some additional assumptions are needed for the argument to work (see Briggs (2009) and Weisberg (2007)). Whether this motivation for Reflection is successful will depend on whether Bayesian conditionalization is a requirement of rationality, and whether the additional assumptions needed for the argument are correct. For these reasons, I will not rely too heavily on this particular motivation for Reflection.

(2) *More Evidence is Better* - Another way to motivate Reflection (and my personal favorite) appeals to the claim that, insofar as we care about having true beliefs, or accurate credences, more evidence is better. In other words, the expected accuracy of our attitudes

---

<sup>13</sup> If you have the view described in footnote 11, this restriction may not be necessary. Also, I don’t mean to be saying that one should never defer to another agent about a *de se* proposition – only that one shouldn’t *always* defer in such cases.

increases with increased evidence. If this is right, then the reason to match your current attitude to your later, more informed attitude is that your later attitude is based on more evidence than your current one and so has a greater expected accuracy.

Why think that the expected accuracy of one's attitudes increases with increased evidence? First, the claim is rather intuitive. To see this note that when the accuracy of your attitudes matters a lot, you will tend to gather more evidence. For example, suppose that you are currently 0.7 confident that it is safe to eat some variety of wild mushroom. But now you encounter such a mushroom and are considering eating it. Suddenly, the accuracy of your attitude matters much more. After all, if you don't eat the mushroom and it's safe, you miss out on a delightful gustatory experience, whereas if you eat the mushroom and it's not safe you die! A reasonable thing to do in this case is to gather more evidence. Why? Presumably because you expect that the accuracy of your attitude will increase with an increase of evidence. Additionally, there are formal proofs given for the claim that (at least on some plausible ways of measuring accuracy), expected accuracy increases with greater information.<sup>14</sup>

(3) *The Disjunction of Evidence Argument* - Jonathan Weisberg (2007) proposes (though does not endorse) the following way of motivating Reflection: When you learn what your future (rational) attitude will be, you are, in effect, learning a disjunction composed of all the different bodies of evidence that rationalize your later attitude. (Note that, since this argument relies on disjunctions of evidence, the argument assumes that evidence is propositional). So, for example, suppose you learn that you will rationally have a 0.9 credence in  $p$  upon receiving some more information. Now consider every potential body of evidence that is a superset of your current body of evidence that would warrant a 0.9 credence in  $p$ . Call these " $E_1 \dots E_n$ ." When you learn that you will later rationally have a 0.9 credence in  $p$ , you learn that:  $E_1$  or  $E_2$  or  $E_3$  or  $\dots E_n$ . Furthermore, assuming that the  $E_i$  form a partition, the fact that each  $E_i$  warrants a .9 credence in  $p$  implies that the disjunction of the  $E_i$  also warrants a 0.9 credence. Since you know the disjunction, you should now have a 0.9 credence in  $p$ .

I hope that I have convinced you that Reflection, or some version of it, is a plausible principle. But Reflection is just a special case of DEFERENCE. To motivate DEFERENCE, we also need it to be appropriate to defer to *other* agents who are more informed than you.<sup>15</sup> I will

---

<sup>14</sup> Paul Horwich (1982) has a nice discussion of this (p.122-129).

<sup>15</sup> For DEFERENCE to hold in full generality it must apply to other agents, future time slices of yourself, and *past* time slices of yourself. Reflection covers future time slices. The arguments here will motivate deferring to both other agents and past time slices but I will be restricting the in-text discussion to other agents for ease of exposition.

do this by showing that the second and third motivations above apply equally well to cases of deference to other rational agents, provided we assume:

**Agent Neutrality:** What it is rational to believe about a (de dicto) proposition given a body of (de dicto) evidence doesn't depend on who is evaluating the evidence.<sup>16</sup>

This is a plausible assumption that has been defended most extensively by evidentialists (people who think that what it is rational to believe depends entirely on one's evidence)<sup>17</sup> so I will not provide a defense of Agent Neutrality here. I will just note that without an assumption like Agent Neutrality, or some restricted version of it, it is hard to make sense of why we *ever* defer to more informed rational agents. After all, even if I know that S is more informed and rational, if I reject Agent Neutrality, I will think that part of the reason that it is rational for S to, say, believe p is the fact that it is S that is evaluating the evidence. But why should I care about what is rational for S to believe, unless that has some bearing on what it is rational for *me* to believe? So, if you reject Agent Neutrality, you incur the burden of explaining why it is that, in many cases, deferring to more informed rational agents is appropriate.

With Agent Neutrality in hand, we can see how the second and third motivations for Reflection motivate the more general principle, DEFERENCE. Let's start with the second motivation for Reflection, the *More is Better* motivation. The thought behind this motivation is that you should judge that the later attitudes you will have, which are based on more information, have greater expected accuracy than your current attitudes. Assuming Agent Neutrality, what it would be rational *for you* to believe given an increased body of evidence, is the same as what it would be rational for some other agent, S, to believe given that evidence. Thus, if you judge that later attitudes of *yours* would have greater expected accuracy, you will also have to judge that some more informed agent, S, will have attitudes with greater expected accuracy than your current ones. After all, we are assuming that S's more informed attitude *is the very same attitude* that you would have were you to have this additional information, so it cannot be that the expected accuracy of one is lower than the expected accuracy of another. So, with the Agent Neutrality assumption, the second motivation for Reflection carries over to DEFERENCE.

Let's now consider the third motivation for Reflection, the *Disjunction of Evidence Argument*. Since we are assuming Agent Neutrality, we know that that what it is rational for S to believe given her larger body of evidence is the same as what it would be rational for

---

<sup>16</sup> If you are a permissivist, you can replace this principle with something like: "What it is rational to believe given a body of evidence *and some set of rational standards of reasoning* (or "*and some prior probability function*") doesn't depend on who is evaluating the evidence."

<sup>17</sup> See, for example, EK in Conee and Feldman (2004) p.29.

you to believe were you to have S's evidence. Now suppose you learn that S rationally has some attitude A towards p. Call the potential bodies of evidence (which are a superset of your own) that would warrant *you* having attitude A towards p "E<sub>1</sub>...E<sub>n</sub>." Since you know that S has attitude A towards p, and that the attitude S has is the very same attitude that it would be rational for you to have, were you to have S's evidence, you learn that E<sub>1</sub> or E<sub>2</sub> or E<sub>3</sub> or...E<sub>n</sub>. Since each of these E<sub>i</sub> warrants attitude A towards p, supposing that they are mutually exclusive, you ought to now adopt attitude A towards p.

In this subsection, I began by providing three motivations for Reflection, which is a special case of DEFERENCE. To defend the more general principle, DEFERENCE, I then argued that the second and third motivations apply just as well to other agents as they do to future time slices of yourself, provided we assume that what it is rational to believe given a body of evidence doesn't depend on who is evaluating the evidence (Agent Neutrality). In the next section I will show why this principle conflicts with CONCILIATIONISM.

### 3. The Argument for a Conflict

In this section, I am going to present two cases which are each meant to illustrate the conflict between CONCILIATIONISM and DEFERENCE. The first case will involve deference to another agent. The second will involve deference to a future time slice.

#### 3.1 The Cases

Case #1:

##### WEATHER PREDICTION

Anna is a meteorologist and is trying to predict whether it will rain on 1/1/2014. (We will call the proposition that it will rain on 1/1/2014 "Rain"). Anna knows that Bob, another meteorologist, has more evidence than she does. (Maybe her evidence concerns weather patterns for 10 years and Bob's evidence concerns weather patterns for 50 years). Anna and Bob also both rationally believe that Bob is under the influence of a drug which is likely to make his weather predictions unreliable. In fact, people under the influence of this drug do no better than chance when making meteorological judgments. Anna learns that, for this reason, after Bob examined the evidence, he suspended judgment on Rain.

CONCILIATIONISM says that it is irrational for Bob to be confident that it will rain. Let's stipulate that, given Bob's evidence about the effects of this drug, suspension of judgment is the rational attitude for him to adopt, and that Anna knows this. DEFERENCE says that it is rational for Anna to defer to Bob (she knows that Bob is more informed and that his attitude is rational). Together they say that it is rational for Anna to suspend judgment. But this is clearly false. If Anna's evidence supports the proposition that it will rain, it would



make no sense for her to suspend judgment simply because the unfortunate soul down the hall has suspended judgment after learning that he might be under the influence of a drug. Since the combination of CONCILIATIONISM and DEFERENCE yield the result that Anna should suspend judgment, but Anna should not suspend judgment, one of these principles must be false. Thus, CONCILIATIONISM and DEFERENCE are incompatible.

I will discuss a possible objection to the judgments I gave in this case in a moment, but first let me first present the second kind of case:

#### Case #2:

##### FUTURE DRUG

At  $t_1$  you have evidence E and suspend judgment on p. At  $t_2$  you gain a bit more evidence which causes you to adopt credence c. You also learn, at  $t_2$ , that, tomorrow, you will gain a lot more evidence about p. However, you will likely be under the influence of a reason distorting drug which disrupts your ability to integrate complex bodies of evidence. You may reliably be able to tell what each piece of evidence supports conditional on E, *individually*, but you will not be able to tell what attitude is supported by your *total* evidence. You learn that, for this reason, you will later suspend judgment on p.

CONCILIATIONISM says that it is rational for your later self to suspend judgment. DEFERENCE says that it is rational for your current self to defer to your later attitude. Together they say that it is rational for your current self to suspend judgment. But this is clearly false. The fact that you will later suspend judgment *because you will be under the influence of a drug* has no bearing on what you should believe now concerning p. For note that, since you won't be able to rely on your ability to evaluate the new evidence you get tomorrow, *whatever evidence you get tomorrow* will result in a suspension of judgment.<sup>18</sup> So the fact that you will suspend judgment tomorrow has absolutely no bearing on whether p. Thus, you should still have the attitude that your current body of evidence supports. Since the combination of CONCILIATIONISM and DEFERENCE yields the result that you should suspend judgment now, but, in fact, you should not suspend judgment now, one of these principles must be false. Thus, CONCILIATIONISM and DEFERENCE are incompatible.

Let me address one worry one might have about my description of these cases. I will state the worry as an objection to WEATHER PREDICTION but the objection applies equally to both cases. You might think that DEFERENCE doesn't yield the result that Anna should defer to Bob, because Anna doesn't know that Bob is *rational*. After all, Bob might be under the influence of a reason distorting drug!

---

<sup>18</sup> This is an instance of what Brian Weatherson (ms.) calls "judgments screening evidence."

Although it is true that Bob, *the person*, may not be rational when under the influence of the drug, the question we are considering is whether or not Bob's *attitude* is rational, and if CONCILIATIONISM is correct, then Bob's attitude *is* rational. Why is the rationality of the *attitude* what matters rather than the rationality of the *person* holding the attitude? If we consider the original motivations for DEFERENCE we will see that the considerations do not motivate a claim about deferring to rational *people*, only to rational *attitudes*. I (and others) have argued that the reason it is advisable to match our attitudes to the attitudes of experts (when these attitudes are rational) is that such attitudes have greater expected accuracy than our current attitudes. Note that if somebody who is very knowledgeable, but also quite irrational, has, in some particular instance, an attitude which you know to be rational towards a proposition that is based on more evidence than your attitude, that attitude *still* has greater expected accuracy than your current attitude. The person's general dispositions to form rational or irrational attitudes do not matter. It is the rationality of the *attitude* that makes the attitude worth deferring to. Thus, even if the fact that Bob might be drugged will lead him to have *some* irrational attitudes, so long as we are conciliationists, and claim that it is rational for Bob to suspend judgment in Rain, DEFERENCE implies that suspension of judgment in Rain is an attitude that Anna should defer to.

It is cases like WEATHER PREDICTION and FUTURE DRUG that illustrate that there is a tension between CONCILIATIONISM and DEFERENCE. I think that this tension reveals a yet deeper tension between two notions of rationality, and it is to this issue that I now turn.<sup>19</sup>

## 4. Two Notions of Epistemic Rationality

### 4.1 Two Desiderata

In the previous section, I argued that there is a tension between two popular views in epistemology: DEFERENCE, according to which we should defer to rational attitudes that are based on more information than our own, and CONCILIATIONISM, according to which agents shouldn't rely on their reasoning when they think that such reasoning will yield results that are no better than chance. This tension, I think, points to a deeper tension between two different "jobs" that we want the notion of epistemic rationality to do for us.

On the one hand, we want the principles of epistemic rationality to be useful in guiding our deliberation. When we deliberate about whether *p*, what we tend to be interested in is whether or not *p* is true, and we are hoping, in deliberation to maximize the expected

---

<sup>19</sup> Lasonen Aarnio (forthcoming) and Sepielli (ms.) also think that higher order evidence considerations motivate a distinction between two kinds of rationality though their notions and motivations are somewhat different from mine. Ichikawa and Jarvis (2013) develop a related distinction which I will discuss later in the paper.

accuracy of our attitudes.<sup>20</sup> For the notion to be guidance giving, in this sense, something *roughly* like the following must hold:

**Guidance Giving Condition (GGC):** For any  $p$ , the attitudes towards  $p$  that result from adopting the principles of epistemic rationality, for the purposes of deliberation concerning  $p$ , must have greater expected accuracy than any alternative attitudes.

You might be wondering why the condition is stated in terms of attitudes that “result from adopting”, rather than attitudes that “are recommended by,” or “in accordance with,” the principles of rationality. The reason is that I do not want principles like “the truth principle” which says: “Have credence 1 in all truths and credence 0 in all falsehoods” to count as guidance giving. It is true that forming attitudes in *accordance* with such a principle will maximize my expected accuracy. However, it is *not* true, or at least not obviously true, that the attitudes that I would end up with as a result of *adopting* this principle.<sup>21</sup>

The second job that we want the notion of epistemic rationality to do for us is help us in our evaluation of the attitudes of other agents. To figure out what rules govern the evaluation of other agents we need to think about why we’re interested in evaluation in the first place. One of the primary purposes of epistemic evaluation is to help us determine when we should defer to the judgments of others.<sup>22</sup> If we know that someone has all the evidence that we have and more, it may be a good idea to defer to them, *provided that the beliefs that they have are rational*. No matter how much evidence someone has, if we don’t think that person’s attitudes are ones that are supported adequately by her evidence, we shouldn’t defer to her.

**Evaluative Condition:** The principles of rationality should be such that it is rational to defer to more informed agents who are conforming to the principles of rationality.

There are, of course, other “jobs” that we might want the notion of rationality to do. Many of these can be done by one of the notions I will be describing (I will list some of these “jobs” later on). However, I’m not committed to the claim that we need *only* two notions to do all of the work that we want rationality to do. I am only saying that we need at least two.

---

<sup>20</sup> Some of the authors who have defended the role of rationality in guiding us towards truth are Cohen (1984), Conee (1992), Kornblith (1993) and Wedgwood (2002).

<sup>21</sup> Lasonen Aarnio (forthcoming) puts this point slightly differently. She says: “such rules don’t offer enough guidance as to what exactly it is that one should do in order to believe the truth.”

<sup>22</sup> Some authors who have defended this sort of evaluative role for epistemic notions are Craig (1990), Reynolds (2002) and Dogramaci (2012).

## 4.2 Why One Notion Can't Satisfy Both Desiderata

The reason that one notion can't satisfy both desiderata is that, as I will argue in a moment, the principles that satisfy the guidance giving condition need to be *agent-relative* and the principles that satisfy the evaluative condition need to be *agent-neutral*. Since no single set of principles can be both agent relative and agent neutral, there is no coherent notion of epistemic rationality that can play both of the roles that we want it to play. First, two quick clarifications concerning my usage of the term "agent relative."

There are two ways in which principles of epistemic rationality might be said to be agent relative. The first way in which principles can be agent relative is for there to be no single set of principles which are the principles of epistemic rationality full stop. There are only principles of rationality that hold relative to particular agents. In other words, *which principles govern the rationality of an agent's attitude depends on facts about the agent*. Alternatively, principles can be agent relative in the sense that, although all agents are bound by the same set of principles, the principles will recommend different beliefs depending on who the agent is. For my purposes, it doesn't much matter which of these we use because any theory that is agent relative in the first sense can be easily translated into a theory that is agent relative in the second sense and vice versa. For this reason, when I talk about the principles of rationality being agent relative, I will sometimes talk as if they are agent relative in the first sense and sometimes as if they are agent relative in the second, depending on which will make exposition easiest. Also, note that in asking whether the principles of rationality are agent relative or agent neutral I am asking whether there are cases in which two agents with the same *de dicto* evidence, ought to have different attitudes towards a *de dicto* proposition (the reason for this restriction are the same as the reasons discussed in section 2.2).

I will now argue that the principles that satisfy the guidance giving condition are agent relative and the principles that satisfy the evaluative condition are agent-neutral.

### *Why The Principles that Satisfy the Guidance Giving Condition are Agent-Relative*

The reason that the principles that satisfy the guidance giving condition will be agent relative is rather simple. Recall that the principles that meet the guidance giving condition for me will be the principles that are such that *trying to follow them* will maximize my expected accuracy. But what attitudes will result from one agent trying to follow a set of principles will differ from the attitudes that will result from another agent trying to follow that very same set of principles.

Consider an analogy. Someone is asking me for directions to the post office in the town in which I live. There are two routes to the post office: one involves going down some main streets and is about a half hour walk. The other involves some shortcuts and twists and turns and is about a twenty minute walk. If the person asking me for directions is

someone who is not familiar with the city and appears to have a terrible sense of direction, I will give them directions that will take them along the longer route. If, however, the person is very familiar with the city, I will give them directions that will take them along the shorter route. The reason for this difference is that, although, if the tourist *followed* my directions for the shorter route, it would take him less time to get to the post office than it would if he followed the longer route, the path that would result from his *attempting* to follow the directions along the shorter route is longer than the path that would result from his attempting to follow the directions that take him along the longer route (since he would likely get lost). What this demonstrates is that when we are trying to guide somebody towards some sort of goal (getting to the post office, having true beliefs), what directions will best achieve this aim depends on facts about the agent.

To make the case for agent relativity more concrete, I will appeal to CONCILIATIONISM. First, I will argue that CONCILIATIONISM is a good guidance giving principle<sup>23</sup> and I will then show (though it might be obvious by this point) that it implies agent relativity. The first task is to show that CONCILIATIONISM is a good guidance giving principle. To do this, I will contrast conciliationism with Thomas Kelly's view, which I'll call "anti-conciliationism" (What I call "anti-conciliationism" is a disjunction of his 2005 and 2011 views, and the argument I will be advancing applies to both disjuncts).

ANTI-CONCILIATIONISM: Suppose that S starts out with a credence  $c_1$  in  $p$ . Then, S receives  $E$  and, on its basis forms a different credence,  $c_2$ , which, *in fact, is the rational response to E*. If, *independently of S's reasoning about p*, S rationally believes that she will do no better than chance at evaluating  $E$ , S's credence in  $p$  should not equal her prior credence in  $p$ . Rather, S's credence should either remain at  $c_2$ , or fall somewhere in the open interval between  $c_1$  and  $c_2$ . (However, if S did not reason correctly on the basis of  $E$ , she should not maintain her new credence).

Which is better from a guidance giving perspective: CONCILIATIONISM or ANTI-CONCILIATIONISM? To figure this out, we need to think about what would result from someone attempting to follow ANTI-CONCILIATIONISM. In fact, I think, it is somewhat unclear what would be involved in trying to follow ANTI-CONCILIATIONISM, and this, in itself, is a strike against its guiding-givingness, but I will consider the two most plausible candidates and show that neither results in greater expected accuracy than CONCILIATIONISM.

Imagine yourself in a situation in which, independently of your reasoning about the issue, you think there is a decent chance that you evaluated your evidence poorly. Perhaps this is a case like DRUG, where you learn that you were under the influence of a reason

---

<sup>23</sup> See Lasonen Aarnio (2010) for an argument in a somewhat similar spirit.

distorting drug when you reasoned about whether  $p$ . Now you are going to try and follow ANTI-CONCILIATIONISM, so you think to yourself: “If, in fact, I reasoned correctly, then my credence should be close to my original answer, but if I did not reason correctly I should revise.” What would be the result of trying to follow this advice? Two options. Option #1: You appeal to your reasoning about  $p$ , examine it and conclude that you reasoned correctly. Therefore you maintain confidence or move your credence somewhat in the direction of your prior credence, but not all the way. Option #2: You don’t appeal to your reasoning about  $p$ . You think: “I have no idea whether I reasoned correctly or not, so I have no idea what the principle recommends.”

We can’t expect great results from you trying to follow the principle in the first way. For despite your best intentions, we know that, in the relevant cases, you will end up near a credence that is *irrational* given the first order evidence a large percentage of the time. So suppose that your attempt to follow anti-conciliationism takes the form of option 2. You simply have no idea, in the relevant cases, what the principle recommends. If that’s right, we, the theorists, can’t make very strong predictions about what you will do. Perhaps you’ll guess, perhaps you’ll pick a credence out of a hat, or perhaps you’ll jump in a lake. Note that I am not assuming that in order to predict that an agent will follow a principle, the agent needs to be *certain* that the conditions in the antecedent are met. I am only claiming that it would be an unreasonable for us to predict that the agent will follow the principle in instances in which *she has no idea* whether the conditions are met. So either, as a result of adopting the principle, the agent’s beliefs will track what she *thinks* the first order evidence supports, or, it’s entirely unclear what she would do as a result of adopting the principle. In either case, it’s plausible that we can expect better results by adopting conciliationism.

I hope to have convinced you that, when it comes to responding to evidence of our own fallibility, CONCILIATIONISM wins the guidance giving contest. But as noted in Christensen (2010), CONCILIATIONISM is an agent relative principle. What an agent ought to believe depends on what limitations the agent rationally believes *she* has. To illustrate, consider:

#### SAME EVIDENCE

Charlie and Dave are each doing a simple logic problem. Their evidence consists of the premises of the problem, and the correct answer (either  $p$  or  $\sim p$ ) is entailed by their evidence. They also both rationally believe that Dave is under the influence of a drug which makes him very unreliable with respect to simple logic problems.

In this case, Charlie should believe either  $p$  or  $\sim p$ , but, according to CONCILIATIONISM, Dave should suspend judgment (or, at very least, he should not be highly confident in  $p$  or  $\sim p$ ). Thus, according to CONCILIATIONISM, there are cases in which agents with the same de

dicto evidence should believe different things concerning a de dicto proposition. CONCILIATIONISM, then, is a good guidance giving principle, but it is agent relative.

In this subsection, I started by offering some general considerations in support of the claim that principles that are good for giving guidance will be agent relative. I then picked a particular principle (CONCILIATIONISM) which scores well on the guidance giving scale, and noted that it is agent relative (in the sense that I am interested in). This demonstrates that the principles that satisfy the guidance giving condition will contain agent relative principles.

### *Why The Principles that Satisfy the Evaluative Condition are Agent Neutral*

I will now argue that if epistemic rationality is to play a role in evaluation for the purposes of determining whether to defer to other agents, the principles of epistemic rationality must be agent neutral. The reason this is so is simple: if I am going to defer to an agent's attitudes, it will only make sense to do so if I know that *her* have a certain attitude, makes it appropriate for *me* to have that attitude. This will only be guaranteed if what attitude the evidence supports doesn't depend on which agent is evaluating it.

Consider an analogy. Eleanor and Frances are going to a wedding. Eleanor is a fashion expert and Frances is not. Since Frances is always apprehensive about what to wear to such occasions she decides that she will defer to Eleanor and wear whatever she is wearing. As it happens, Eleanor is the bride and is going to be wearing a long lacy wedding gown. Frances decides to rent a long lacy wedding gown to wear to Eleanor's wedding.

Clearly, Frances' policy of deferring to Eleanor in this way is not going to end well. This is because what is appropriate to wear to a wedding *is not* agent neutral. What is appropriate to wear depends on *who you are*. This is why matching your choices to that of someone else, even if that person is an expert on appropriate wear to weddings, is not advisable. The same is true when we are considering matching our belief states to those of other agents. This is only a sensible policy if what belief is appropriate doesn't depend on who the agent is.

### **4.3 The Two Notions and the Solution**

In the previous subsection, I argued that the principles that satisfy the guidance giving condition are agent relative and the principles that satisfy the evaluative condition are agent neutral. What we need, then, are two distinct notions of epistemic rationality. Since no single set of principles can do both of the jobs that we want epistemic rationality to do for us, we should designate one notion as the one that satisfies the guidance giving condition, and another as the one that satisfies the evaluative condition. I call the notion that satisfies the guidance giving condition "what one ought to believe." This notion is useful from the first personal, deliberative perspective and is sensitive to agent limitations.

I call the notion that satisfies the evaluative condition “what the evidence supports.” This notion is useful from the third personal, evaluative perspective and is not sensitive to agent limitations. (Nothing much hangs on the names that I have chosen, but I chose them as I did because I think that the notion that satisfies the guidance giving condition is the deontic notion, the one that governs how agents *ought* to form beliefs).

Although I have introduced the notion of evidential support as the notion that satisfies the evaluative condition, Jonathan Jenkins Ichikawa and Benjamin Jarvis (2013) argue that a notion like this one plays other important roles as well. For example, it is the notion that we use to make sense of epistemic improvement and it is also the notion that we use to explain the success of rational agents. I am in complete agreement with them, and I don’t think that the evaluative role I described is privileged in any way. I simply focused on the evaluative condition and the guidance giving condition since these roles enable to us to see clearly what is going on in the cases of higher order evidence. There is much more to be said about both of these notions and the roles they each play than I can say here.

I will now proceed to show how this distinction can help deal with the puzzling cases I described in section three. Since we now have two distinct notions of epistemic rationality, we can ask: which of the following versions of CONCILIATIONISM is correct?<sup>24</sup>

CONCILIATIONISM<sub>O</sub>: What ***one ought to believe*** depends, in part, on what limitations the agent ought to believe she has. In particular, if, *independently of S’s reasoning about p*, S ought to believe that she will do no better than chance at evaluating some body of evidence, S’s credence in p ought to equal her prior credence in p.

CONCILIATIONISM<sub>ES</sub>: What ***the evidence supports*** depends, in part, on what limitations the evidence supports the agent having. In particular, if, *independently of S’s reasoning about p*, the evidence supports belief in the proposition that she that she will do no better than chance at evaluating some body of evidence, the credence supported by S’s evidence is S’s prior credence.

Since, as I argued earlier, CONCILIATIONISM is a good guidance giving principle, it should be clear that it, at very least, is a principle that governs what one ought to believe. Now, of course, sometimes (indeed, much of the time) principles that govern what one ought to believe will also be principles of evidential support, so the fact that CONCILIATIONISM is a principle of what one ought to believe doesn’t itself show that it’s *not* a principle of

---

<sup>24</sup> I should note that rationality-like notions show up a number of times in the principle, and so there are other versions of the principle, which involve mixing and matching. However, these are the two versions that will be relevant to our discussion.



evidential support. However, since, as noted earlier, CONCILIATIONISM is an agent-relative principle, it cannot be a principle of evidential support. So CONCILIATIONISM<sub>0</sub> is true but CONCILIATIONISM<sub>ES</sub> is false.<sup>25</sup> Now let's consider two versions of DEFERENCE.

DEFERENCE<sub>0</sub>: If you know that S has strictly more de dicto evidence than you do (that is, S has all the de dicto evidence that you have, plus more), S has no less self-locating information than you do, and that S's doxastic attitude towards a de dicto proposition p is ***an attitude that S ought to have*** then you ought to adopt S's doxastic attitude towards p.

DEFERENCE<sub>ES</sub>: If you know that S has strictly more de dicto evidence than you do (that is, S has all the de dicto evidence that you have, plus more), S has no less self-locating information than you, and that S's doxastic attitude towards a de dicto proposition p is ***supported by the evidence*** then you ought to adopt S's doxastic attitude towards p.

Since the notion of evidential support tracks attitudes that are deference-worthy and DEFERENCE just is the principle that tells you what attitudes to defer to, it should be clear that DEFERENCE<sub>ES</sub> is correct. But what about DEFERENCE<sub>0</sub>? I have already argued that what an agent ought to believe is an agent relative notion, and that, in general, if we're trying to figure out which attitudes are deference-worthy we want an agent-neutral notion (recall the case of Eleanor and Frances at the wedding). So a version of DEFERENCE that told you to defer to attitudes that agents ought to have, would sometimes yield the wrong results.

Having established that CONCILIATIONISM<sub>ES</sub> and DEFERENCE<sub>0</sub> are false, we are in a position to return to the puzzling cases in section 3. Recall that the question that arose in those cases was the following: why should we not defer to more informed attitudes in cases in which those more informed attitudes are based on a conciliatory response? Since CONCILIATIONISM<sub>ES</sub> is false, the more informed attitudes, in both the weather case and the future drug case, are attitudes the agents ought to have, but are *not* attitudes that the evidence supports. The right version of deference only tells you to defer to the attitudes of more informed agents when those attitudes are ones that the evidence supports. Since this condition isn't met in the problematic cases, there is no puzzle as to why one shouldn't defer.

Of course, just because we shouldn't *always* defer to more informed attitudes that agents ought to have, doesn't mean that we never should. It is reasonable to defer to an

---

<sup>25</sup> Christensen (2010) has written that agents who respond in a conciliatory way to higher order evidence are embodying a certain kind of "epistemic imperfection." The claim that conciliatory responses are not attitudes that the evidence supports can be seen as a way of putting one's finger on the kind of imperfection such agents are embodying.

agent who is believing as she ought to if the attitude that results from *her* believing as she ought to is more expectedly accurate than the attitude that results from *you* believing as you ought to. This could happen in two different ways. First, it might make sense to defer if you know that, neither you, nor your friend, has the attitude supported by the evidence, but that your friend has more evidence than you do. It could also happen if you have the same evidence, neither of you has the attitude supported by that evidence, but your friend is doing better than you are at evaluating the evidence (though she's not doing perfectly).

Life is complicated. Unfortunately, there is no general formula one can give concerning when we should defer to an agent who is believing as she ought to. The important point is that an agent believing as she ought to is not *sufficient* for deference. I have been arguing that, for the purposes of deference, we need an agent neutral notion. But this agent neutral notion need not be an all or nothing notion. Just as we can rank agents on the basis of how much evidence they have, we can rank agents on the basis of how well they are evaluating the evidence *with respect to the agent neutral notion*. If they are doing better than we are then, even though they are not doing perfectly, deference might still be reasonable.

I have argued in this section that we need two notions of rationality: one to for the purposes of guiding deliberation and one for the purpose of evaluating other agents. After distinguishing these two notions I argued that we should only accept the “what one ought to believe” version of CONCILIATIONISM and the “what the evidence supports” version of DEFERENCE. These principles don't yield conflicting results in the cases described in section three.

## 5. Some Loose Ends

Recall that, in section 2, I gave arguments for CONCILIATIONISM and DEFERENCE. In the previous section, I argued that CONCILIATIONISM<sub>ES</sub> is false and that DEFERENCE<sub>0</sub> is false. What I want to do now is show that the original motivations for CONCILIATIONISM do not motivate CONCILIATIONISM<sub>ES</sub> and the original motivations for DEFERENCE do not motivate DEFERENCE<sub>0</sub>.

### 5.1. Why the Original Motivations Don't Motivate CONCILIATIONISM<sub>ES</sub>

Throughout the paper, two motivations were given for CONCILIATIONISM. The first was that such a principle is necessary to explain our intuitive judgments in cases like DRUG. The second was an argument given in section three, aimed at showing that attitudes that result from adopting a conciliatory principle have greater expected accuracy than attitudes that result from adopting alternative principles. With respect to the first motivation, the intuitive judgments are, I think, sufficiently explained by CONCILIATIONISM<sub>0</sub>. So we don't need CONCILIATIONISM<sub>ES</sub> to explain such judgments. The second argument appealed directly to the thought that adopting conciliatory policies will increase the expected accuracy of our

attitudes. But such considerations only motivate the claim that we *ought* to be conciliatory, since such considerations appeal to the importance of rationality being guidance giving. As I argued earlier, we cannot expect the notion that satisfies the evaluative condition to be the same as the notion that is guidance giving. So when we argue for a principle by appealing to its' guidance-givingness, we are arguing for a principle about what one ought to believe, and not about what the evidence supports.<sup>26</sup>

## 5.2 Why The Original Motivations Don't Motivate DEFERENCE<sub>0</sub>

Recall that there were three arguments given for Reflection (and two of these, were then used to motivate the more general principle, DEFERENCE).

- (1) Reflection follows from Conditionalization
- (2) More Evidence is Better
- (3) The Disjunction of Evidence Argument

I will consider each separately.

(1) The first motivation, stemming from Conditionalization, does not motivate DEFERENCE<sub>0</sub>. This is because, as David Christensen (2010) has pointed out, in a case in which you are responding to higher order evidence in the way CONCILIATIONISM recommends, you are not conditionalizing. Let's use the future drug case as an illustration of this. In the future drug case, CONCILIATIONISM tells you to suspend judgment, regardless of what particular evidence you will get in the future. But before you learned about the drug, your conditional credences in *p* given such evidence, *and the fact that you will be drugged in the future* do vary based on what the evidence is. So, by suspending judgment, you are not updating by conditionalization. Furthermore, CONCILIATIONISM frequently requires less than full confidence in propositions entailed by your evidence, which is also inconsistent with conditionalization. Since CONCILIATIONISM is true of what you ought to believe, conditionalization cannot be a requirement of what you ought to believe, and so cannot be used to motivate DEFERENCE<sub>0</sub>.

(2) The second motivation for DEFERENCE discussed earlier involved the thought that, the more evidence you have, the greater the expected accuracy of your attitudes. This, in general, is correct. However, in cases in which you respond to higher order evidence in the way CONCILIATIONISM recommends, you won't judge the expected accuracy of your attitudes to be greater than the expected accuracy of your current attitudes. This is because the additional evidence that you will be getting is no help to you if you are not in a position to judge what that evidence supports. Since what the more informed agent's attitude is, in these cases, doesn't depend on the actual evidence the agent gets, the attitude formed in

---

<sup>26</sup> This is not to say that principles of evidential support cannot also be guidance giving. I am only saying that scoring well with respect to guidance giving cannot, in and of itself, motivate a principle about evidential support.

response doesn't get the benefit that additional evidence usually brings with it. Recall that we only regard attitudes based on more evidence as more expectedly accurate when we know those attitudes involve an appropriate response to the evidence. What we've seen is that there are two different ways of understanding "appropriate response." It can be understood as "a response the agent ought to have" or "a response that the agent's evidence supports." Since the response the agent ought to have will not actually benefit from the additional evidence in the kinds of cases we're discussing, the sense of "appropriate" that needs to be plugged in to the "more evidence is better" slogan needs to be evidential support. Also note that the formal arguments given for the claim that attitudes based on more information have greater accuracy than your current attitudes assume that you will be conditionalizing, which, as I showed above, is not the case when you are responding in the conciliatory manner to higher order evidence.

You might be wondering the following: if, in fact, you don't judge that your later attitude in the future drug case will be more accurate than your current attitude, why don't you stick to your current attitude even once you get the additional evidence? What is the purpose of changing your credence if you don't think the resulting credence will be more accurate? The answer is as follows: although you *now* regard your current attitude as having greater expected accuracy than your later attitude, once you get the additional evidence, you will *not* regard your earlier attitude as more expectedly accurate.

(3) The final motivation for DEFERENCE was the disjunction of evidence argument. The thought was that, if you know what your later attitude will be, you know that the disjunction of all the bodies of evidence warranting that attitude is true. Since all of the disjuncts warrant your future attitude, and the disjuncts are mutually exclusive, the disjunction as a whole warrants your future attitude. Return again to the future drug case. Note that the fact that, in the future, you will suspend judgment *does not* tell you that the evidence you will get will be of some variety that warrants suspension of judgment. For, once again, suspension of judgment is the attitude you will adopt *regardless* of what evidence you receive. So when you know that you will respond in a conciliatory way to higher order evidence in the future, you are not learning the disjunction of the kind that is needed to warrant your adoption of your future attitude now.

In the previous section I argued that CONCILIATIONISM<sub>ES</sub> and DEFERENCE<sub>0</sub> are false. In this section I argued that the original motivations for CONCILIATIONISM and DEFERENCE do not motivate the false versions of these principles. Fortunately, as I argued in section three, it is only the false (and unmotivated) versions of the principles that are in tension with one another. So all is well.

## 6. Conclusion

I have argued that a single notion of “rationality” cannot serve both as the notion we use to guide deliberation and as the notion we use to evaluate others. The notion we use to deliberate with should account for the possibility of our limitations, since deliberative principles that take our limitations into account give us a better shot at having accurate credences. Believing as we ought to is doing the best we can with what we’ve got. The notion that we use to evaluate with, however, should be agent-neutral, and thus should not account for agents’ cognitive limitations. This is because when we’re trying to figure out who we should defer to, we don’t want two people to be rated equally just because they’re both “doing the best they can with what they’ve got.” What we need then, are two notions: one to deliberate with, and one with which to evaluate others.

## References

Arntzenius, F. (2003) “Some Problems for Conditionalization and Reflection.” *Journal of Philosophy* 100(7): 356-371.

Briggs, R. (2009) “Distorted Reflection.” *Philosophical Review* 118(1): 59-85.

Christensen, D. (2007) “Epistemology of Disagreement – The Good News.” *Philosophical Review* 116(2) :187-217.

Christensen, D. (2010) “Higher Order Evidence.” *Philosophy and Phenomenological Research* 81(1): 185 – 215.

Conee, E. (1992) “The Truth Connection.” *Philosophy and Phenomenological Research* 52(3): 657-669.

Conee, E. and Feldman, R. (2004) *Evidentialism*. Oxford University Press.

Cohen, S. (1984) “Justification and Truth”, *Philosophical Studies* 46(3): 279-295.

Craig, Edward (1990). *Knowledge and the State of Nature*. Oxford University Press.

Dogramaci, Sinan. (2012) “Reverse Engineering Epistemic Rationality” *Philosophy and Phenomenological Research*, 84(3): 513-530.

Elga, A. (2000) “Self-locating belief and the Sleeping Beauty problem” *Analysis* 60(2): 143-147.

Elga, A. (2007) “Reflection and Disagreement” *Noûs* 41(3): 478-502.

Elga, A. (2010) “How to Disagree about How to Disagree” In Ted Warfield & Richard Feldman (eds.), *Disagreement*. Oxford University Press.

- Elga, A. (ms.) "Lucky to Be Rational."
- Horgan, T. (2004) "Sleeping Beauty Awakened: new odds at the dawn of a new day." *Analysis* 64(1): 10-21
- Horowitz, S. and Sliwa, P. (ms.) "Respecting *All* The Evidence."
- Horwich, P. (1982) *Probability and Evidence*. Cambridge University Press.
- Ichikawa, J.J. and Jarvis, B.W. (2013) *The Rules of Thought*. Oxford University Press.
- Kelly, T. (2005) "The Epistemic Significance of Disagreement" in T. Szabo Gendler and J. Hawthorne eds. *Oxford Studies in Epistemology Vol. 1*. Oxford University Press.
- Kelly, T. (2011). "Peer Disagreement and Higher Order Evidence." In Alvin Goldman and Dennis Whitcomb (eds.) *Social Epistemology: Essential Readings*. Oxford University Press.
- Kornblith, H. (1993) "Epistemic Normativity", *Synthese* 94(3): 357-376.
- Lasonen-Aarnio, M. (2010). "Unreasonable Knowledge." *Philosophical Perspectives* 24(1): 1-21.
- Lasonen-Aarnio, M. (forthcoming). "Higher Order Evidence and Limits of Defeat." *Philosophy and Phenomenological Research*.
- Perry, J. (1977) "Frege on Demonstratives" *Philosophical Review* 86(4): 474-97.
- Reynolds, S. (2002) "Testimony, Knowledge, and Epistemic Goals." *Philosophical Studies* 110(2): 139-161.
- Sepielli, A. (ms.) "Evidence, Reasonableness and Disagreement."
- van Fraassen, B. (1984) "Belief and the Will." *Journal of Philosophy* 81(5): 235-256.
- Vavova, K. (ms.) "Irrelevant influences."
- Weatherson, B. (ms.) "Do Judgments Screen Evidence?"
- Wedgwood, R. (2002) "The Aim of Belief", *Philosophical Perspectives*, 16: 267-297.
- Weisberg, J (2007). "Conditionalization, Reflection, and Self Knowledge." *Philosophical Studies* 135(2): 179-197
- White, R. (2010) "You just Believe that Because..." *Philosophical Perspectives* 24(1): 573-615.