

# The Epistemology of Disagreement

Justin Blank

February 24, 2010

How should we respond to disagreement with people we regard as our peers?<sup>1</sup> It is an important feature of our intellectual lives that we take such disagreements seriously, that we seek out alternate perspectives on our ideas and compare our reactions to evidence with other investigators. We undeniably benefit from subjecting our views to others' scrutiny, and it is no accident that the majority of cranks and crackpots work outside of any community of serious investigators. Yet there still remains a question of how it is that disagreement makes a difference. When we disagree with someone and they provide us with what we recognize as a good argument that we have reasoned incorrectly, this gives us reason to revise our views. This fact leaves the question of how we should react when someone disagrees with us and we are not moved to accept some argument that they offer. In that case, should we take the disagreement itself as a reason to revise our views?<sup>2</sup>

One extreme is the position that disagreement itself gives us no reason

---

<sup>1</sup>Theories of disagreement differ on how they define peerhood. I will not adopt any particular definition, as nothing in my discussion requires it.

<sup>2</sup>In what follows I use 'disagreement itself' or 'fact of disagreement' where I need to stress that it's mere disagreement that gives me reason to change my views, though I'll continue to use 'disagreement' where the context is clear.

whatsoever to alter our views. When I disagree with a peer, if I reasoned correctly, I should maintain my current opinion, while if I reasoned incorrectly I ought to revise my opinion, but in light of my original evidence, not the fact of disagreement. At the other extreme is the view that each of us should split the difference between our views, regardless of which one of us is right.

I believe that there is a simple taxonomy of disagreements which will allow us to show their epistemic significance. In constructing the taxonomy, I consider how we react to the evidence in isolation—without the influence of comparing our answer with other individuals—and ask how our reaction to the evidence may change over time. In some cases, our views will be stable and will represent our best response to the evidence, or our considered opinion, while in others, our answers will only represent a temporary resting place in an ongoing process of wrestling with our evidence, even after careful reflection. Much of the literature on disagreement abstracts from our beliefs' potential for change over time, but I will argue that it is a crucial element in understanding disagreement.

After laying out the taxonomy, I will consider several cases, showing how they intuitively fit with the taxonomy, and how it illuminates how we should respond to disagreement in them. The view that I develop will be a variant of Tom Kelly's total evidence view, and I accept most of what he says concerning that view. However, while his view is relatively noncommittal on how disagreement adds to our evidence, my discussion allows me to clarify the underlying factors that make disagreement important in particular cases. Moreover, I present an argument for the total evidence view which does

not depend on some of the controversial assumptions of his bootstrapping arguments.

## 1 The one where I talk about the taxonomy

I claimed that much of the literature on disagreement works with an overly simplified picture of how we respond to evidence. Typical descriptions involve us considering the evidence, then forming a particular level of credence in the proposition we're evaluating, with no mention of either intermediate inferences or change over time. On such a simplistic picture, it is as if our credence in a proposition is simply triggered by our evidence, so that we are passive and lack any insight into the process. I doubt that any of the authors involved would consciously endorse such a picture of reasoning, but their discussions and arguments seem to reflect this picture. What I wish to focus on is the process by which someone responds to their evidence in isolation, prior to consulting other people.

The range of topics where the simplified picture is apt is quite small.<sup>3</sup> Sometimes we are in possession of an effective procedure for producing an answer. Ordinary arithmetical calculations are such a case. In other cases, we lack any explicit procedure, but still move directly from evidence to a conclusion without mediating reasoning. My judgments about whether it will rain work this way. I have neither special meteorological knowledge nor the experience a lifelong farmer might have. I just look at the sky and

---

<sup>3</sup>Strictly, whether the simplified picture applies or not is a product not of the topic, but the topic given a particular reasoner's knowledge and strategies for approaching the topic, but I will suppress this complication.

make a guess, which will have limited reliability. Nor can I typically improve my judgment by spending more time looking at the clouds and pondering whether they mean rain. I simply don't know enough about the relevant features—that dark clouds often accompany rain is about the limit of my knowledge. A third type of case occurs where we have a reliable ability to make judgements about a topic, but lack awareness of how we do so, as do the proverbial chicken sexers, who reliably judge the sex of baby chicks but do not know how they do so. Such cases differ from my judgment of the weather because the individuals are reliable and know that fact. Fourth, there are simple perceptual judgments. These judgments are reliable and require no complicated inferences. Moreover, we ordinarily have introspective access to the grounds of our judgment. Consider a spectator judging which horse wins a race or whether a tennis ball lands within the lines of the court. Sometimes they might draw on explicit knowledge of perspective, optical illusions, etc., to make a judgment, but assume that our spectator has no knowledge that might be relevant to the race he's watching. Then, his only rational option is to believe what he seemed to see, while acknowledging the possibility that he misperceived the situation.

With these four types of cases set aside, we can see that many topics do not fit the simplified picture. None of the conclusions in this paper depend on the exact range of cases that fit the simplified picture.<sup>4</sup> Still, the importance of noting this simplification is enhanced by the fact that *none* of

---

<sup>4</sup>However, it is at least dialectically significant that some authors hew closely to the simple cases when motivating the equal weight view. Elga, for instance, initially begins by mentioning a judgment concerning the deterrent effect of the death penalty, but switches to a horse race when it comes time to make an argument from bootstrapping (Elga 487).

the cases which initially drove the disagreement literature—religious, ethical, scientific or philosophical disagreement—seem to fit the simplified picture.<sup>5</sup>

What distinguishes the simple cases is that there is a narrowly circumscribed range of evidence that the agent is capable of calling upon.<sup>6</sup> Of course, within the simple cases there are further divisions. In the mathematical cases, the agent knows an explicit procedure which guarantees the correct answer if appropriately followed. In the simple perceptual case, the observer has clearly delineated evidence which has a clear bearing on the subject and no further evidence. Perhaps the same is true of the chicken sexing case, though there the chicken sexer cannot specify what his evidence is. In the weather case, the observer may know many facts that would be evidence, but he can neither identify them as evidence nor make a reliable judgment in response to them. In spite of those differences, in each of the simple cases the agents will take into account all the evidence to which they are capable of responding.<sup>7</sup> Moreover, their judgment on the basis of that evidence is likely to be stable regardless of how much they think about the situation.

In more complicated cases we have a wider range of evidence and have to select the pieces of evidence that are most important for us to focus on and weigh them against each other. We may not have rules to guide us in this

---

<sup>5</sup>On some views of what ethical reasoning involves, it will bottom out in judgments which fit the picture, with disastrous consequences. See Setiya

<sup>6</sup>“Calling upon” is deliberately vague. Is there an epistemological term that’s intentionally ambiguous between how you respond to evidence you know you have and evidence that you don’t know is evidence?

<sup>7</sup>See the section on stability for elaboration of what kind of capability I mean. Perhaps on some ordinary understanding of capability, the agents are ignoring evidence that they are capable of responding to, if so, then I am considering a restricted notion of capability.

process, and when we do, they will usually include *ceteris paribus* clauses, or terms whose application requires non-mechanical judgment. However, this lack of unambiguously applicable rules does not mean that we have no grip on which pieces of evidence matter. Otherwise, the complicated case would be like the layman's judgment of the weather. In complicated cases, if we consider a particular factor, we will usually know whether or not it is relevant. We will even typically recognize whether that piece of evidence is defeated, if we attend to the relevant defeater. What is difficult is recognizing all the pieces of evidence we possess that might defeat a given piece of evidence, and judging how heavily to weigh that evidence. We can improve our understanding of how any particular piece of evidence matters by reflecting hard on what its defeaters might be, or constructing arguments that it is particularly important.<sup>8</sup> Yet in light of the non-mechanical judgment previously alluded to and sheer quantity of evidence, our apprehension of the relevant evidential relations will almost always be partial. That is not to say that we cannot have justified beliefs. Either the weight of the evidence may strongly favor one answer, or we may reason well about a difficult case. If we reason well, we select the most important pieces of evidence for consideration, and we note the most important potential defeaters. Our verdict about the case reflects each of those subsidiary judgments, and they serve to justify our actual belief.

For concreteness, suppose we are considering which of two basketball

---

<sup>8</sup>I can't remember this for sure, but I think I may be using defeater in a non-standard way here or throughout the paper. I seem to recall a distinction between undermining defeaters, those which cause a fact to no longer play its role as evidence, and defeaters which merely outweigh some piece of evidence. If that's the actual distinction, I think I'm only using 'defeaters' to refer to undermining defeaters.

players is better. There are no definitive rules for moving from statistical information about the players to relative assessments of quality, but we still have warrant for valuing their points per game or shooting percentages.<sup>9</sup> What complicates our considerations is how the factors interact, and the range of possible defeaters each one has. If one player has slightly more points per game than another, we still need to consider their shooting percentages. But shooting percentage has a complex relation to number of shots taken, percentage of shots that are 3-pointers, a player's position, and the quality of his teammates, among other factors. The list of relevant factors is practically unlimited, and if we had to consider all possible defeaters before rendering a judgment, we would be epistemically paralyzed. Still, there are easy cases: all stars will typically exceed ordinary players in so many categories that considering a few of the more salient measures will suffice. Even in closer cases, we may be able to construct good, if defeasible, arguments that one player is better than another.

We may apply many different strategies, but they all involve taking subsidiary questions which we could consider at length and attempting to pass over them with simple arguments or summary judgments. Doing so exposes us to a risk that our summary judgment will be less accurate than had we considered the issue in detail. Consider how often in constructing a philosophical argument we pass over a particular case, assuming that it presents no additional complications, only to be proven wrong. Although this is always a possibility when we make judgments about complex cases,

---

<sup>9</sup>A measure like PER aims to be a rule, but in practice, we recognize its limitations. Depending on our level of knowledge, it might be reasonable to let PER screen the rest of our evidence, but we might also know enough to take it as one factor among many.

it does not preclude our having justified beliefs.

What separates the simple cases from the complex ones is that in the simple cases, our beliefs are stable. I will call them “stable under reflection”, since the relevant aspect of their stability is how they change when we reflect on our evidence for them. My beliefs about the chance of rain may change, or I could forget which horse seemed to win, but these are non-rational changes that are not based in my reasoning from the evidence I possess. As demonstrated by the four sorts of simple cases, stability under reflection is not directly produced by the quality of our evidence or the accuracy of our response to it.<sup>10</sup> However, it is the quality that matters for our response to disagreement.

## 2 Stability

Stability under reflection is a modal notion. In calling my belief stable, we are not considering how my belief will change in the future with further reflection, as I may never consider the topic after comparing beliefs with my interlocutor, but how my beliefs might change. It’s necessary to clarify the modality involved. If we accept a weak reading of this modality, then our taxonomy will treat too many of our methods of forming beliefs as equivalent—it will view them as highly subject to change. At the other extreme, a strategy for answering a question with results that are highly subject to chance may still only produce one answer in actual circumstances,

---

<sup>10</sup>Though the following conjecture seems promising: in simple cases we have extremely circumscribed evidence: either clearly definitive evidence or little evidence whatsoever. Complex cases involve more evidence of a more diffuse kind.

no matter how much time I spend. Such a reading of the modal terms will treat all of our strategies as on a par—highly resistant to change.

What we need is an intermediate modal status. I will not try to precisely specify this modality in standard modal vocabulary, so much as point to its essential features. First, it is a normative modality insofar as the changes of belief must meet a standard of reasonableness. Not just any change of belief will count—it must be a change based on a response to evidence. My beliefs about the weather may change after I have initially looked at the sky and predicted that it would rain. I may look at the same clouds and think “now I don’t feel like it’s going to rain.” Although the change in opinion may change a true belief to a false one or vice versa, the transition is neither epistemically praiseworthy or blameworthy.<sup>11</sup>

In general, I am incapable of improving my warrant for my belief by thinking about the weather, hence my belief is stable. That said, the change involved need not be produced by an impeccable or fully explicit chain of reasoning to count against a belief’s stability. If I become more confident that the recession will end because I attend to a change in employment figures, that will be sufficiently reasonable even if I have overlooked evidence that suggests that the change is merely seasonal. In such a case, there are available defeaters for the inference, but a reasonable judge could fail to note them. As a deliberate simplification, we can say that the change in belief should consist of an inference that is *prima facie* reliable or valid.

---

<sup>11</sup>This statement excludes the rare case where my initial belief was obviously defective, even to me. If I look at an utterly clear sky and somehow judge that it will rain then I can recognize that my original judgment was unreasonable. The cases I’m concerned with occur when my initial belief was rational, or alternately, I had no way of knowing it was irrational.

Second, stability is normative insofar as what matters is not whether there are close counterfactual scenarios where I do in fact change my position on the situation. Psychological research and common experience both give us reasons to believe that once we have accepted a position, we behave as “intuitive lawyers” seeking out confirmatory evidence and ignoring inconvenient evidence. In that regard, adjusting our beliefs about controversial topics is not our natural behavior. If we just consider close counterfactual scenarios, then we would give undue weight to the beliefs of the most stubborn judges, counting their beliefs as stable under reflection. So stability under reflection must somehow incorporate how we *should* respond to our evidence.

Third, the notion of stability under reflection must limit the possible inferences to those we could plausibly make. Consider my judgments about the economy. Although my own economic knowledge is limited, many aspects of economic models are quite simple, relying on little math beyond high school algebra.<sup>12</sup> Some models obviously are motivated by or involved parameters derived from econometric research that I have no access to. Still, other models involve no such element—though they have empirical assumptions, much of the insight they offer does not require specific economic data.<sup>13</sup> Should we count my beliefs about the economy to be quite unstable because I have the mathematical skills to construct a model which undermines all the ar-

---

<sup>12</sup>By model, I do not mean the type of thing which would be used to make quantitative predictions about the course of the economy in the coming year but rather models similar to the IS-LM model or more up to date analogs thereof.

<sup>13</sup>To make the point that follows even stronger, we could note if I’ve been reading the newspaper in full, I may know an awful lot of measures of what’s going on in the economy, in case we think that this sort of knowledge is necessary for developing the models that we’re discussing.

guments that I've previously encountered?<sup>14</sup> I think not. These models were constructed and tested over a long period of time, and often required substantial insight by major economists to be developed, even if their mathematical underpinnings are quite simple. If we wish to explain the disconnect between mathematical simplicity and substantial advance in understanding, we might note that the success of the models is based on their creators' ability to identify the key phenomena and isolate them from a great deal of noise. However much I think about economics alone in my study, I am unlikely to repeat that feat, but this is not because the mathematics of the model are beyond me.

To some extent, the second and third constraints on stability under reflection pull in opposite directions. Oversimplifying, we could say that the second constraint tells us to abstract away from the individual's cognitive limitations, while the third constraint requires us to take them into account. Yet the constraints are compatible. Consider how we evaluate a student's paper on a complex subject. While our comments might note abstruse points that the student has missed, our evaluation of the students' work will depend on a conception of which insights are feasible for someone with his level of exposure to the material, time constraints, etc.<sup>15</sup> Stability under reflection is similar to a more demanding analogue of the evaluation that we place upon students' work. A bad paper will be so unstable that the student

---

<sup>14</sup>See Krugman's "There's Something About Macro" <http://web.mit.edu/krugman/www/islm.html> for an argument that such toy models play a role in policy debates.

<sup>15</sup>Cherniak 1986 argues that we need a conception of 'minimal rationality' for understanding the conditions that an agent must satisfy in order to be a believer, and it is plausible to think that such minimal standards extend to other types of evaluation.

would see errors upon a first rereading, while an ideal paper would be the best that student could produce until a later stage of their philosophical education.

Although the above three constraints do not define the concept of stability under reflection, I believe that they give us a sufficient grip on it to characterize our beliefs about various areas as more or less stable under reflection. Even if the idea of stability under reflection is somewhat indeterminate, it will do the work for which it is needed. In what follows, I will argue that stability under reflection is the primary determinant of whether disagreement should lead us to change our beliefs. In particular, it is only when we face disagreement involving beliefs that are stable under reflection that disagreement itself makes an important contribution to what we should believe. When our belief is unstable under reflection, disagreement at most serves to direct our attention to aspects of our epistemic position that were already accessible to us.

### **3 Disagreement**

How should one respond to disagreement in the simple cases? In the three non-mathematical simple cases, you should split the difference between your view and your peer's view. The best you can do is regard your opinion as a probabilistic indicator of the proposition you're judging. Even if you haven't explicitly thought of the matter that way, you should do so, given what you can say about how you reached your opinion. Your judgment that your interlocutor is a peer commits you to taking his opinion as a probabilistic

indicator of equal reliability, so you should give his opinion the same weight as your own view.<sup>16</sup> Thus far, our verdicts about cases match both the total evidence view and Kelly's equal weight view.<sup>17</sup>

In a case where my belief is unstable under reflection, that is because my reaction to the evidence depends on how well I grasp the relevant evidentiary relations. In responding to the evidence, I do not merely say what strikes me as right, or what I seem to have seen. Rather, I make inferences about what the evidence supports.<sup>18</sup> A naive thought is that if I am right, this precludes splitting the difference, because doing so would make my belief about the case inconsistent with my beliefs about the weight of the evidence I possess. Since those beliefs about how the evidence bears on the case at hand are correct, I should take them to be my guide and maintain my initial response. The proponent of the equal weight view has obvious rejoinders to the naive argument. For it seems that the naive argument works only if I assume that I'm the one who responded correctly to the case at hand, and in a case of disagreement, I'm not entitled to that assumption. Given

---

<sup>16</sup>On some definitions of a peer, this may not follow, however those definitions of peerhood make the Equal Weight View less plausible, so I will not discuss them.

<sup>17</sup>I'm leaving aside the case of arithmetic calculation, because it presents difficulties. On the hand, we have definitive evidence that indicates a particular answer via an effective procedure. On the other, we seem to think it is equally likely that we are wrong as that our interlocutor is. It is difficult to know how to model the process of reasoning in terms of beliefs or credences, a fact that is related to the problem of logical omniscience.

<sup>18</sup>I may not form beliefs about evidence *per se*, but surrogate beliefs such as those about probabilities will do just fine. Even more ordinary, idiomatic ways of expressing arguments will do: if I think that given a fact in my evidence, something "must be" or "probably is" true, that thought plays much the same role as an explicit belief about my evidence. It does matter that these colloquial expressions are less precise, and how we represent evidential relations will matter to how much higher-order evidence we receive from disagreement, as I will discuss below. Still, the argument that we cannot split the difference will go through just as well regardless of how we represent these evidential relations.

that I can't just assume that I'm right, I should split the difference about the case. Then what the inconsistency between my updated credence and my background beliefs shows is that I should also split the difference about those background beliefs.<sup>19</sup> In what follows, I will demonstrate that the inconsistency between my background beliefs and my judgment about the case cannot be eliminated by changing those background beliefs. Although it requires elaboration to see this fact, the naive thought is correct.

The case is easiest to see if we make a pair of simplifying assumptions. First, suppose that in determining my opinion about the topic, I come to a set of explicit beliefs about how the evidence bears on the subject, and second, that given those background beliefs about my evidence, the conclusion deductively follows.<sup>20</sup> While these assumptions give a simplified picture of reasoning, they do fit some actual cases. One strategy we may employ in simplifying a complex topic is to attempt to construct a deductive argument, even if the underlying evidence is inductive.

If my background beliefs suffice to make my conclusion the result of a valid inference, then splitting the difference will not be an option. For so long as I am confident in the premises of the inference, I should be confident of the conclusion.<sup>21</sup> If I give up on the conclusion without changing my background

---

<sup>19</sup>By background beliefs, I mean the beliefs that support my conclusion, whether they be things I believed before thinking about the topic, or whether they be mediate conclusions from my initial evidence. Thus 'background beliefs' does not fit perfectly.

<sup>20</sup>If we change this from "deductively follows", to "follows according to some algorithm" we get the calculation case discussed earlier, with the assumption that all of the intermediate steps in the calculation are matters about which I have explicit beliefs. This assumption is a substantial idealization, but I believe it lends some weight to thinking that we should not split the difference in cases of calculation.

<sup>21</sup>Of course, the question of exactly how credences should transfer from a set of premises to the conclusion of an argument from those premises may be complicated. Here I am assuming that the premises reach whatever threshold is necessary for me to be quite

beliefs, I will be committed to beliefs that are obviously inconsistent.

The consequence is that to split the difference, I must revise the premises of my argument in such a way that they are no longer inconsistent with my conclusion. However, since my belief is unstable, I will typically have further reasons for the premises of my argument. To give up on those premises, I must alter my opinion of how those beliefs support my premises, and this process can be iterated, as I then examine my evidence for those background beliefs. What the result should be will depend on whether my response to my initial evidence was correct or not. If my initial response was incorrect and my interlocutor responded correctly, then there should be no difficulty.<sup>22</sup> In revising my premises, I will eventually come to one of my faulty assumptions. Perhaps I will, upon revisiting the matter, see that I should not have made the assumption. Revising it will introduce no tension into my beliefs.<sup>23</sup> The case where one faulty assumption is all that separates my position from my interlocutor's will be somewhat rare, but it is not impossible. We certainly do make arguments based on assumptions that we would not accept if we considered them in the right light. For whatever reason, we may not attend to these assumptions until confrontation with an interlocutor makes them salient, at which time we easily see that they are ungrounded.

If, on the other hand, I reasoned correctly, and my interlocutor merely  

---

confident in the conclusion.

<sup>22</sup>I'm describing the cases as if there is always one correct and one incorrect party, with the correct one wholly grasping the right answer. Relaxing this assumption would complicate the discussion, but I believe it would change nothing fundamental.

<sup>23</sup>Strictly, this is false. If the disagreement rested solely on my having made that assumption, then upon giving it up, I should not split the difference, but simply adopt my interlocutor's position. What is accurate is that relaxing the false assumption will remove a source of tension in my beliefs, without adding a new source of tension.

had one faulty assumption, it is difficult to see how I can revise my beliefs. For my situation is just the reverse what it would be in the case where I reasoned poorly. To split the difference, I would have to introduce an inconsistency into my beliefs, corresponding to the faulty assumption of my interlocutor.<sup>24</sup>

It remains to extend the argument to less idealized cases and deal with an objection. Our idealization involved assuming that after considering our evidence we summarized its import by giving a deductive argument. When we do not form such a deductive argument, the problem with our beliefs will not be as obvious as before.<sup>25</sup> Still, in this case, my beliefs fail to cohere with each other. Certain pieces of evidence seem to me to be highly relevant, others seem irrelevant. I have some grasp on which pieces of *prima facie* evidence are defeated and which are not. By splitting the difference, I fail to base my opinion on the pieces of evidence which I take to be most relevant.

From my presentation so far, it might appear that I think disagreement should provide no reason to alter our beliefs in complicated cases. If so, my argument would seem to prove too much, for it is hard to accept that absolutely nothing should change when we are exposed to peer disagreement. However, I need not maintain that disagreement is idle in complicated cases.

How I represent these facts about relevance will depend on my understanding of the topic, and this will alter how I can respond to disagreement. What is at issue is how my understanding of my evidence constrains my

---

<sup>24</sup>I'm here assuming that my splitting the difference will involve adjusting my background beliefs to match my interlocutor's beliefs as closely as is possible. See the later discussion of what portion of their reasoning parties to the disagreement must share.

<sup>25</sup>Non-deductive arguments can typically be transformed into deductive ones, but this often requires accepting trivializing or question begging premises.

belief. There are multiple dimensions on which my understanding can vary. At a minimum, they can differ in how well I understand the import of the evidence I'm considering, as well as how strongly that evidence constrains the proper answer. In some cases, our evidence consists of *purely contributory reasons*. Purely contributory reasons would be evidence that consists of quite a number of facts, each of which makes an independent contribution to the proper credence, some positive, some negative, so that the proper credence is just a result of summing the reasons.<sup>26</sup> In such a case, disagreement might be quite important. For all the reasons for believing  $p$  that I see, there might be countervailing reasons against  $p$  that I haven't attended to. If I could literally enumerate each piece of evidence and assign it a weight, disagreement would be unimportant. Since I cannot, I should take my interlocutor's belief as a good reason to reduce my credence in  $p$ . I can reason as follows: "my interlocutor has the same evidence as I do. I noticed more reasons for believing  $p$  than not  $p$ , but he's a peer, and he judged not  $p$ . He wouldn't have done that without attending to reasons that not  $p$ , but that suggests that maybe there's more of them than I thought." Since nothing that I'm aware of rules out that the reasons for believing not  $p$  are more numerous than I'd thought, I can easily reduce my credence in  $p$ . However, the possibility of splitting the difference depends on how well I can represent the balance of reasons. If it were the case that I simply had a hunch that the

---

<sup>26</sup>Upon reflection, it's not clear how many examples fit this mould—at least I have had trouble coming up with any that in areas I know about. The illustration that comes to mind is a list of pros and cons for some decision, but upon reflection, it seems quite rare that such a list would actually consist of independent reasons, each of which can seem to make some additive contribution to the desirability (or not) of the course of action. Perhaps an example like judging the potential severity of an impending flu season might do.

reasons for  $p$  were greater than the reasons for not- $p$ , then I will be capable of splitting the difference. For my hunch will be much like the beliefs that appeared in simple cases that were stable under reflection. The more that I can construct arguments from verifiable premises towards the conclusion that more reasons favor  $p$ , the less I will be able to split the difference.<sup>27</sup>

This is a good time to mention an issue raised by my earlier discussion. My own argument claimed that we often cannot split the difference because it would make our beliefs inconsistent but just now, I allowed that you might split the difference in a case characterized by contributory reasons. However, if my initial reaction was correct, my beliefs will be just as inconsistent. For the actual balance of reasons will favor  $p$ , and there will be relatively few reasons to believe not- $p$ . The defender of the equal weight view should say that this is ok, and what the case shows is that certain kinds of inconsistency are not especially bad. Now, inconsistency might be acceptable: there are reasons to doubt that any process which produces inconsistent beliefs is thereby made unacceptable. We almost inevitably have inconsistent beliefs, and although that fact does not prove that consistency is not an ideal of rationality, it is suggestive that many such inconsistencies need not impair our ability to answer the questions we set for ourselves.<sup>28</sup> More to the point, cases like the preface paradox have been used to argue that there is

---

<sup>27</sup>Note that this fits with the taxonomy of stable vs. unstable beliefs. If it just seems that there are stronger reasons for  $p$ , then that belief will be stable. If there is a structured inference to the claim that stronger reasons support not- $p$ , then it will be unstable.

<sup>28</sup>Inconsistent beliefs may be functionally segregated so that they are not applied for reasoning about the same topics, or it may be that inconsistent beliefs persist because we do not consider them at the same time, though we could easily resolve the inconsistency. See Cherniak 1986 for detailed discussion of these claims.

no rational obligation to make our beliefs consistent.<sup>29</sup> However, whatever force these qualifications might have against a categorical requirement of consistency, they do not affect the point at hand. The kind of inconsistency created by splitting the difference concerns a single topic, involves explicit judgments that the subject can recollect, and involves giving up the conclusion of an argument that the subject previously accepted. It is hard to see what standards of reasoning could remain if we allowed that such an inconsistency was acceptable.<sup>30</sup> On the other hand, the inconsistency in the purely contributory case is mere preface style inconsistency.

Returning to the ways my understanding of the evidence can vary, there are other cases where disagreement may matter. In the purely contributory case, the weight of each individual reason was easily apprehended, and the difficulty arose in weighing them against each other. At the other extreme, there are cases where the considerations I base my judgment on are extremely unclear. If I consider a mathematical question where it is not clear what proof strategies are viable (take the continuum hypothesis for example), but where I initially think that there is a reason to favor one position. If I consult with other similarly knowledgeable individuals, many of whom disagree, I should at least lose confidence that my answer correct, if it's not the case that I should come to agree with their stance.<sup>31</sup> In such

---

<sup>29</sup>Or so I've heard. I'm not sure if I've encountered the argument, though I've been told Christensen 2004 gives it.

<sup>30</sup>Nor is the mention of "standards of reasoning" an indication that this argument depends on the existence of rational requirements not reducible to facts about reasons, such as those denied by Kolodny.

<sup>31</sup>At the extreme, such a case might collapse into the weather case—one's opinion about unsolved conjectures would merely have a probability of being true or false. However, there are forms of expert judgment that do not look like this, where we have a somewhat inarticulate grasp of reasons for believing or disbelieving a particular position. Consider a

a case, my first order evidence may end up swamped by what I receive from disagreement.<sup>32</sup>

A second way of seeing the problem with splitting the difference, which we can take either as a separate argument, or an illustration of what is so bad about the inconsistency highlighted in the first argument, concerns our ability to defend or adjust our revised credence. Suppose that after conferring with my first peer, and splitting the difference with him, I come into contact with a second peer, with whom I then disagree. Leave aside the question of how much weight I should give my current view, relative to his, and ask how I can defend it to him. I can give my reasons for my initial answer, prior to splitting the difference, but these will not match my updated credence, but rather my credence before splitting the difference with my first peer. Here we run into an ambiguity involved in the literature on disagreement: I know of no accepted characterization of what peers must share in cases where they split the difference.<sup>33</sup> That is, what portion of our reasoning on the disputed matter do we share with each other before splitting the difference? One possibility is that we do not share our intermediate stages of reasoning with each other. We simply view the evidence, form credences, and then split the difference.<sup>34</sup> In this case, it is clear that after

---

philosopher's judgment about whether a particular argumentative strategy has a chance of succeeding, prior to developing that argument.

<sup>32</sup>Note that this is also a case of a stable belief, and one where we incline towards splitting the difference.

<sup>33</sup>(Elga 484) seems to suggest that parties to the dispute might share nothing but their resulting credences, (Christensen 5) seems to imagine a case where two individuals each recognize that their respective theories each succeed on one of two conflicting epistemic desiderata. (Feldman 419) clearly imagines that the two individuals have fully shared their reasons and arguments.

<sup>34</sup>Note how this fits with the simple picture.

splitting the difference, I will be unable to fully defend my credence to a third party. I will be able to explain the rationale for my contribution to my credence, but not my original peers' rationale.

On the other hand, if the peers wholly articulate their reasoning for their positions, then the prior discussion of stability under reflection suggests that their reasoning will at least partially change as they take into account each others' arguments. Even if they continue to take opposed positions, they may moderate their credences in virtue of newly recognizing pieces of evidence which are relevant. In this case, the original response of splitting the difference is not directly applicable. If I acknowledge that one of my assumptions was wrong, and raise my credence, then it would be irrational to split the difference based on my starting credence. For that reason, it is difficult to make sense of the equal weight view in contexts where dialogue has the potential to change my mind. The view only has obvious applicability after dialogue's potential has been exhausted.<sup>35</sup>

Alternatively, suppose that I split the difference with my peer, then acquire new evidence. Specifically, say that I gave a high credence to  $p$  which I then reduced as a result of disagreement. Let  $q$  be some portion of my initial evidence which I judged to be important support for  $p$ , and suppose I learn  $r$ , which would defeat the evidential support  $q$  provides for  $p$ . How should I weigh this evidence? No good options are available. If I ignore  $r$ , I am behaving irrationally, so far as I ignore evidence that I take to be relevant, yet if I further reduce my belief in  $p$ , I am at risk of double-counting.  $r$  may defeat  $p$  in virtue of evidential relations which my

---

<sup>35</sup>I'm really not sure how to think about the issues raised in this paragraph.

interlocutor had already taken into account.<sup>36</sup> Since I split the difference with him, my reduced credence in  $p$  already reflects the weight of the relevant evidence.

In fact, the acquisition of new evidence is not even necessary for the point. Since my belief is unstable under reflection, I have an epistemic reason to improve it, at least when that belief becomes salient. My disagreement with my peer means that we are disposed to differ in which inferences we find compelling. So far as I split the difference, I have acted as if I found his inferences compelling, even though I do not. Since I do not find them compelling, it's unlikely I could recognize all the inferences that make sense from his point of view. Therefore, when I update my beliefs, I will most likely fail to update in a way that respects my decision to split the difference. To the extent that I so fail, my beliefs will either return to what they would have been had I never split the difference, or become increasingly inconsistent.

## 4 Higher Order Evidence

The strongest objection to my view would be to highlight how disagreement provides higher order evidence, and to claim that higher order evidence can disrupt one's grip on even valid inferences. So the fact that complying with the equal weight view would render our beliefs contradictory is no objection to the equal weight view, since this often happens as a result of acquiring

---

<sup>36</sup>This can be true despite  $r$  being novel evidence. Perhaps  $r$  was relevant to  $p$  in virtue of justifying an inference to some further fact, but our initial evidence  $E$  also justified an inference to that same conclusion. Still, I did not initially notice that inference. Another possibility is that  $q$  could have separate defeaters, each of which were sufficient to wholly undermine its support for  $p$ .

higher order evidence that we are reasoning poorly. Consider Christensen's take on higher order evidence: "But [higher order evidence] seems, at first blush, to be evidence of a peculiar sort. For one thing, its evidential import is often agent-relative. For another, respecting it can apparently force an agent to fall short in certain ways, by having beliefs that fail to respect logic or basic inductive support relations." (Christensen 10)

Against this objection, we can note that to the extent that discussions of higher order evidence rely on disagreement as a paradigm case which helps motivate the impact of higher order evidence, it would be circular to use conclusions about higher order evidence in order to support a conclusion about how we should respond to disagreement. Still, there are examples other than disagreement that motivate reflection on higher order evidence, and Christensen gives three. However, each of these examples differs from a case of disagreement in an important way.<sup>37</sup> Christensen's cases are a medical resident who has been awake for 36 hours and a case where I've been administered a drug that interferes with my ability to do logic puzzles. In each of these cases, I receive evidence not only that I might be reasoning poorly about the specific case, but also that I have a disposition to reason poorly about a range of cases.<sup>38</sup> As such, if I believe the evidence of my own error, I cannot rectify the situation by attending to my evidence about the

---

<sup>37</sup>I will omit Christensen's final example, where I is attached to a machine that makes me believe I'm in mental state S iff I am not in S. The difficulty is that it's unclear how the scenario works. Is the mental state one that I'm supposed to be aware of via introspection? Does the machine operate on S or my judgment about S, and how many milliseconds does it take to alter S or my judgment? There are too many unanswered questions for the example to be useful.

<sup>38</sup>Christensen notes that the drug and sleep deprivation cases are more similar to each other than to the disagreement case, though he obviously doesn't endorse the line of thought I'm presenting here.

original topic. My beliefs are stably bad.<sup>39</sup> As such, the case differs from a case of disagreement where my beliefs are unstable under reflection. Facing disagreement, I can continually improve my epistemic position by considering my evidence, but I cannot do this if I end up adopting contradictory beliefs as a result of following the equal weight view. In short, we should adopt the following principle: I should only accept contradictory beliefs on the basis of higher order evidence if what that evidence shows is that I cannot avoid contradictory beliefs. This principle is imprecise, because we have not clarified when we *can* avoid contradictory beliefs, and a full discussion of higher order evidence would have to clarify that notion. However, as it stands, the higher order evidence given by disagreement does not meet that standard under any construal of what we can avoid.

Therefore we can take issue with a concession Tom Kelly makes. He suggests that “Even if one has in fact responded to the evidence impeccably on a given occasion, one might still have reason to doubt that ones performance was impeccable. Such a reason is provided when a peer responds to that same evidence differently.” (Kelly 28)<sup>40</sup> From our discussion of complex cases, we can see this is misleading. It is not that disagreement gives us a reason to revise beliefs which were formed impeccably. It is that it highlights and makes salient ways in which our beliefs were not formed impeccably, by giving us reason to reexamine the arguments that we’ve offered, but also reason to recognize that our grounds for belief were incomplete. The cases

---

<sup>39</sup>We can also imagine a sleep-deprivation case where I merely need to think more before issuing a judgment, but such a case is less interesting.

<sup>40</sup>His use of ‘impeccably’ aside, Kelly seems to recognize this point: “Typically, when one responds reasonably to a body of evidence, one is not utterly blind to the fact that one has done so; on the other hand, such facts are not perfectly transparent either.”

are simply too complicated to simultaneously apprehend everything that matters.

Several topics remain to be addressed. Although I did not frame my analysis of disagreement as a more general treatment of higher order evidence, I suspect that with due care, I could show that it fits a particular approach to that topic. Furthermore, it would be ideal to integrate discussion of cases of calculation, but this might require consideration of the difficult problem of logical omniscience.

## References

- [Cherniak 1986] *Minimal Rationality*.
- [Christensen 2004] Christensen, David. *Putting Logic in Its Place* (2004), Oxford.
- [Christensen 2007] pistemology of Disagreement.
- [Christensen ms ] “Higher Order Evidence.”
- [Elga 2006] “Reflection and Disagreement.”
- [Feldman ] “Epistemological Puzzles About Disagreement” in S. Heatherington, ed *Epistemology Futures*. Oxford, 2006.
- [Krugman ] “There’s Something About Macro”. Available at <http://web.mit.edu/krugman/www/islm.html>.
- [Setiya ms] “Does Moral Theory Corrupt the Youth?”