



MONASH University

How Special is Disagreement?

Maxwell Josef Riess

Master of Philosophy

A thesis submitted for the degree of Master of Philosophy

Monash University in 2016

School of Philosophical, Historical & International Studies

Copyright notice

© Maxwell Riess (2016).

I certify that I have made all reasonable efforts to secure copyright permissions for third-party content included in this thesis and have not knowingly added copyright content to my work without the owner's permission.

Abstract

A few years ago philosophers started discussing disagreement. The main question they were seeking to answer was: when I find out someone disagrees with me, what should I do? This discussion has since become the Epistemology of Disagreement. This flurry of interest suggests that there is something philosophically interesting or unique about instances of disagreement, something that the previous decades of epistemology don't quite address. In this thesis I start to examine this idea, that there is something epistemologically unique about cases of disagreement, by analysing some of the theories proposed by philosophers so far. I find that as yet, we do not have reasons to consider disagreement unique and that more could be gained from identifying what is common to the Epistemology of Disagreement and other topics in epistemology.

Declaration

All Students should reproduce this section in their thesis verbatim

This thesis contains no material which has been accepted for the award of any other degree or diploma at any university or equivalent institution and that, to the best of my knowledge and belief, this thesis contains no material previously published or written by another person, except where due reference is made in the text of the thesis.

Signature:



Print Name: ...Maxwell Riess.....

Date: ...14/9/2016.....

Acknowledgements

This thesis is dedicated the following people and organisations, without whom, it would never had been possible.

To the Australian Government, thank you for your generous support in the pursuit of this work.

To Professor Graham Oppy, your patience and generosity are matched only by your thoughtfulness. You put the 'super' in 'epistemic superior'!

To my family for giving me the chance to try.

To Sarah Braybrooke, my perfect partner.

How Special is Disagreement?

By Maxwell Riess

Table of Contents:

Preface	7
1. Introduction to the Epistemology of Disagreement	8
2. What is the right account of Epistemic Peerhood?	28
3. Is Disagreement special?	43
4. Disagreement and 'Belief Independence'	65
5. Concluding remarks	83
References:	87

Preface

This thesis was started in March of 2012 and was largely completed by March of 2014. At that time – for a variety of reasons – I wasn’t able to complete and submit the final thesis, but I was able to come back to it and finalise it in 2016. Because of its history, the thesis is grounded in the pre-2014 literature.

1. Introduction to the Epistemology of Disagreement

In recent years a lively discussion has sprung from philosophers thinking about disagreement. The question most directly discussed in the epistemology of disagreement is: what should I rationally believe when I learn that someone disagrees with me?

In what follows I will explain why disagreement has become an area of interest to philosophers. I will then canvass some of the issues discussed within the debate before then mentioning some of the areas which I think deserve further investigation.

Two Types of Disagreement

Much of the time disagreements arise because different people have different information. As an illustration, here are two contrasting examples of ‘non-peer disagreement’:

Non-peer Disagreement

When meeting with my supervisor recently we began discussing the work of philosopher-X. I was explaining how surprised I was that philosopher-X believes P. At which point my supervisor chimed in and said ‘I think you’re wrong about that, philosopher-X doesn’t believe P’. Now, since they are a far more knowledgeable philosopher than I am, and the work of philosopher-X is one of their areas of specialisation, upon hearing my supervisor’s opinion I immediately changed my view. I no longer think that philosopher-X believes P and my supervisor has taught me something new. Generally speaking, when you learn your teachers disagree with you, the right response is to adopt their perspective. This is part of the lot of being a student.

In contrast, that afternoon when I was teaching an undergraduate course, one of my students disagreed with me that argument-Y is an example of modus tollens; instead they asserted that argument-Y is an example of modus ponens. This time I was unswayed by their belief. They are my philosophical inferior; I know more about the structure of arguments than they do and have thought about it more. This is reflected in my steadfast commitment to my belief in the face of my disagreeing student. Correcting confused students is just part of the lot of being a teacher.

These disagreements are typical in the sense that they arise out of the disagreeing parties having different information. My supervisor knows a lot more; my student knows a lot less. This information imbalance is not only critical in understanding what *causes* the disagreements, but it also plays the principal role in *determining how those involved should*

respond to the disagreement. The fact that my supervisor knows a lot more about philosophy than I do means that when we disagree it is reasonable for me to adopt his view. But my student knows a lot less about philosophy than I do and so when I am the teacher, it seems reasonable for me to resist their opinion and hold on to my original view.

Obviously this type of reasoning doesn't just apply to philosophy teachers and their students but to all cases where one disputant has some epistemic advantage over the other. Parents correct young children, a major experiment with a large and well-drawn sample corrects a smaller observational study, and primary witnesses hold sway over secondary reports. In situations like these, the better-informed party holds court not because they are right necessarily but because they seem to be better informed.

Most people agree that in situations where two parties hold differing opinions on some issue and one of them is far more expert than the other, it is prudent to side with the expert. We can use this consensus to develop a useful 'rule of thumb' on how to respond to similar cases:

When people disagree, the better-informed opinion should hold sway.

This is pretty good first stab at developing a method for resolving disagreements. It clearly is far from complete; we still need to say how we should compare agent views for 'informedness'. This is a problem that we will come back to^{***}, but still this seems like a good place to start.

This rule of thumb has some intuitive appeal. It is simplifying – as a teacher or as a student, the way I should respond to disagreement is relatively clear. Furthermore, my resultant belief is supported by the majority of the disagreeing parties' collective information/credentials, and this seems like a good feature of the rule.

Peer Disagreement

However, things become more difficult when I discover that I disagree with my peer. Unlike my supervisor and my student, my philosophical peer, I think, knows as much about philosophy as I do: it would be false to say that they are either superior or inferior to me in credentials or philosophical status¹. So, later when I am chatting in the postgraduate

¹ Here, philosophical credentials are being used to illustrate the more general notion of 'epistemic status'. Now precisely what factors contribute to someone's epistemic status is a substantial question that I will deal with in detail in a later chapter.

study room and I mention that I think 'Two-Boxing' is the best strategy in Newcomb's problem, I am surprised to discover my philosophical peer disagrees. She tells me that she doesn't think that 'Two-Boxing' is the best strategy in Newcomb's Problem; the puzzle is how should I respond to this 'peer disagreement'²?

Peer disagreements like this one, where neither party is epistemically superior / inferior to the other, have become the focal point in the epistemology of disagreement. This is because cases of peer disagreements appear to be problem cases; saying definitively what to do in this case is more difficult than in the previous two cases. The naive 'rule of thumb' for responding to disagreements, that we should side with whichever opinion is better credentialed, will not help us decide whether we should adopt or resist the views of peers who disagree with us. By stipulation neither of us is philosophically superior to the other. Also, without this relationship between the bearers of the conflicting beliefs in question, the radical epistemic moves – of completely adopting or resisting disagreeing views – seem inappropriate.

Responding to Peer Disagreements: All or nothing beliefs vs credences

So, in the disagreement I have with my peer over Newcomb's problem, if adopting or dismissing their view is inappropriate, then what other options are there? An attractive option is to walk the middle path between the two extremes.

There are potentially two ways of doing this. To see the first of these, notice that when I adopt the belief of my epistemic superior I am arguably performing not one but two epistemic actions. In responding to the disagreement I not only gain the belief $\neg Q$, that 'not-[philosopher-X believes P]' but I also abandon my initial belief Q that 'philosopher-X believes P'. This is clearly desirable since it prevents me from believing contradictory propositions.

Similarly when resisting the belief of my student I am both denying their belief and maintaining my own. In both the situations I end up believing one of the propositions involved and not believing the other. So perhaps by combining these two responses we can come to a way of responding to peer disagreement which is 'in-between' the superior-to-inferior type disagreements. As already mentioned, if a disagreement is between peers who respectively avow the beliefs P and $\neg P$, then responding by believing both seems

² I've borrowed this example from Kelly (2005, 2011).

obviously problematic, as it would require disagreeing peers to believe a contradiction. However, combining the *dis-believing* options is viable. So perhaps when disagreeing with an epistemic peer I should dis-believe both of the propositions initially believed by those involved. Or to put it another way, I should respond to peer disagreements by suspending judgement on the area of disagreement.

This way of responding to disagreement reflects an 'all-or-nothing' view of belief. Within this understanding of belief, for all propositions under consideration one either: believes the proposition, disbelieves the proposition or suspends judgement on the proposition.

But there are a few reasons why an all-or-nothing conception of belief is ill suited to thinking about responding to disagreement. It is typically thought that disagreements influence our beliefs by presenting us with evidence of some kind or other, and evidence is thought to influence our beliefs in a graded way. Evidence for a proposition might be very strong or quite weak; either way it seems reasonable for responses to evidence to reflect that evidence's strength.

This is particularly striking when one considers how to respond to disagreements with numerous peers. Consider my peer disagreement over Newcomb's Problem: as things stand, if one adopts the all or nothing view of belief, when I discover my lone peer in the postgraduate room disagrees with me, I should suspend judgement on the proposition that 'Two-Boxing is the best strategy in Newcomb's Problem'. However, how should I react if I find two peers who independently disagree with me on the best strategy in Newcomb's problem? If this extra peer is to have any additional impact on my beliefs about the disputed issue, then my only option is to now adopt their belief 'that Two-Boxing is *not* the best strategy in Newcomb's Problem'. This may still seem all right, however strangeness appears when we consider how I should react to three disagreeing peers, or 10, or 200. Once I have adopted the belief of my dissenting peers, I have run out of doxastic options, and these extra peers can have no additional effect on my beliefs about two-boxing. However, it does seem that disagreeing with two peers presents one with very different evidence than disagreeing with 200 peers.

There is another way to find a middle path in-between the extreme non-peer responses of the opening examples. If being a philosophical superior allowed me to be unmoved by a disagreement and being a philosophical inferior required me to completely discard my previous opinion, then perhaps in peer disagreements I should change my view, but only partially. So if my philosophical peer disagrees with me over the correct strategy in Newcomb's Problem, then I should become less confident in my previous belief and

become more confident in the belief of my disagreeing peer: both partially adopting and partially resisting their view.

In contrast to the all-or-nothing view of belief mentioned previously, this way of responding to disagreement takes belief as a matter of degree (or credence). Thinking of belief in this way has a number of advantages. Firstly it allows disagreements to act as evidence that can influence our beliefs via degrees, so responding to a disagreement involving two peers will likely be very different from a disagreement involving many peers. Secondly, taking beliefs as credences allows us to capture more instances of disagreement because it will admit cases where, although agents both avow a belief in P, in fact one party is far more confident of P's truth than the other. We could imagine two people who both believe that it is more likely than not that there is life on other planets and because of this they may both evince a belief in extra-terrestrial life. However this still leaves a lot of room for one party to be near certain that there are aliens out there and for the other to be far less confident of the existence of aliens.

Thinking of beliefs as credences not only brings to the fore the possibility of such a disagreement but it also gives us a method for responding to the disagreement which is sufficiently fine-grained to bear out the relevant intricacies of that disagreement.

Conformism and Non-Conformism

Surveying the literature on disagreement, commentators have characterized the debate by looking at how various accounts say we should respond to instances of peer disagreement. These responses are thought to come in two general camps: conciliationist³ (also called conformist) views and steadfast⁴ (or non-conformist) views. Speaking loosely, conciliationist accounts are those that think peer disagreements require substantial belief revision while steadfast views deny that such substantial revision is required in the face of peer disagreements.

In an attempt to map the territory of the current debate quite a lot has been made of these categories and their contrasting take on responding to peer disagreement. However, these categories aren't that useful for two reasons: firstly, because no one in the debate thinks that all cases of peer disagreement require the same response in all circumstances, and secondly, even if there is agreement among how to respond to different cases of peer

³ Authors who are considered to hold conciliationist views are Adam Elga (2007), David Christensen (2007) and Richard Feldman (2009).

⁴ Authors who are considered to hold steadfast views are: Thomas Kelly (2010) and Ernest Sosa (forthcoming).

disagreement, the account justifying those responses are importantly distinct. Grouping accounts of disagreement together on the basis of how they advise agents to respond to disagreement is a superficial taxonomy that doesn't really speak to the content of the views themselves⁵.

There is actually a surprising amount of agreement over the notion that in many instances of peer disagreement the rational response is for those involved to lower their confidence in their initial opinion and raise their confidence in the views of those they disagree with. However this idea does not apply to all instances of peer disagreements discussed in the literature. There are, I think, two types of examples often discussed that elicit opposing intuitions on how rational agents would respond to them.

Two Basic Types of Peer Disagreement

Consider the following two different examples as representatives of two basic types of cases of peer disagreement often discussed:

BILL-SPLITTING⁶: You and I go out for lunch and when the bill arrives we both divvy up what we owe by doing a bit of mental arithmetic. We can both clearly see the bill's total and we've previously agreed to split the bill's total evenly and add a 10% tip. We have a long history splitting bills like this and, although we agree most of the time, on occasions where we disagree, in the past we are each correct around half the time. It turns out that this is one of those times and when we announce the result of our mental calculations, we disagree on what our shares are. I think we each owe \$43 while you think we each owe \$45.

MATH: You and I are classmates revising for a math class by marking a practice exam. We have been classmates for a while now and we often revise together. In fact it's a particularly good arrangement because we are pretty much equally good at math. However it turns out that we have different answers to one of the questions; in answer to the question 'what is $7763 \times 5 = ?$ ', you have given an answer that doesn't end in a 0 or 5, whereas my answer does.

Most commentators broadly agree that these two cases of peer disagreement rationally require different responses from those involved. Most commentators think that in BILL-

⁵ In a series of papers David Christensen (2007, 2009, 2011) has argued that there is a single principle in the debate that divides the conformists from the non-conformist: this is the principle he calls 'independence'. I discuss this principle in detail in a later chapter.

⁶ This example is due to Christensen (2007).

SPLITTING I should lower my confidence in our shares being \$43 and raise my confidence in our shares being \$45. However in MATH it seems I am rational in maintaining my belief that the correct answer ends in either a 0 or 5, and again most commentators working in the area recognise this.

These two examples are emblematic of the two types of cases which are often brought up when discussing peer disagreement⁷. The two cases are similar because, by design, at least up until the moment of disagreement those involved are considered epistemic peers. But there are two pertinent differences between the types of cases that BILL-SPLITTING and MATH represent. These are: the confidence that I have in my initial opinion prior to discovering the disagreement, which is far higher in MATH than it is in BILL-SPLITTING, and my confidence that my disagreeing peer is correct both before and after the disagreement is discovered, which is far lower in MATH than it is in BILL-SPLITTING.

These two factors can be added together to form a handy measure, which I'll call a 'confidence gap'. A confidence gap is the measure of the distance between the confidence someone has in their initial view being correct and the confidence someone has in the view of a disagreeing peer. The widest confidence gaps occur in instances of disagreement like MATH where my confidence in my initial view being correct is extremely high and the confidence I have in my peer's belief being correct is extremely low⁸. I am very confident that I haven't misread the question or your answer (although I am not infallible in this regard), and I am well aware that all multiples of 5 must end in either a 0 or a 5. By contrast the confidence gap in BILL-SPLITTING is quite narrow; it was a mental calculation after all, and I know from experience that it is easy to make mistakes when performing even simple sums in one's head. In addition, your answer in BILL-SPLITTING appears quite plausible and it could easily be correct.

When discussing idealized cases of peer disagreement, we can use the 'confidence gap' to categorize cases as either MATH-like or BILLSPLITTING-like. And as mentioned previously, there is broad agreement over how agents should respond to these different types of peer disagreement. In cases of narrow confidence gaps the consensus intuition is that that peers should move to the middle ground, and in cases of wide confidence gaps

⁷ Oppy (2010) makes a similar 'two type' distinction when discussing a number of the cases often talked about in discussion over disagreement.

⁸ There are arguably further cases of even wider confidence gaps. These are cases where one's belief is 'immediately grounded' via experience or logic. For example if you have a headache and someone thinks you are only feigning illness, then their disagreement shouldn't sway you in the slightest. This would especially be the case if you thought that such beliefs were incorrigible. (Sosa, forthcoming)

the consensus intuition suggests that disagreeing peers may reasonably maintain their positions.

Christensen (2007, 2011) makes a similar distinction, however he distinguishes MATH-type cases from BILL-SPLITTING-type cases only on the basis of one's very high initial credence in the MATH-type cases. For Christensen what distinguishes MATH-type cases where one can legitimately resist the view of a disagreeing peer from BILL-SPLITTING-types cases, where one must accommodate their belief to some significant degree is the fact that in MATH-type cases one's initial confidence in one's view prior to the disagreement arising is far higher than in BILL-SPLITTING-type cases.

However I think this way of splitting the cases is slightly off the mark. While I do have a very high initial credence in MATH (and many other cases like it), there are other cases that evoke similar intuitions where one's initial view doesn't have an ultra high credence. For instance, in another version of the BILL-SPLITTING case, I may not be very confident of my mental arithmetic and consequently I will not be very confident that our shares are in fact \$43. However if you report that you think our shares are \$485 each, then I don't think I need to significantly alter my initial belief. That is because your evinced opinion is so far fetched that it has dramatically widened the confidence gap between our views, and it does this despite the fact that I don't have a particularly high credence in my own view being correct. In other words, by adjusting the confidence gap, a BILL-SPLITTING-type example can become much more like a MATH-type example⁹.

So if there isn't direct disagreement over whether epistemic peers need to change their beliefs in the face of particular disagreements, what exactly is the disagreement over disagreement? The answer to this question is two-fold. Firstly there are different views on which features of a disagreement influence what one is rationally required to believe when responding to disagreement. And secondly, (often because of this first issue) there is divergence over the extent of the belief revision when it is required.

Two Contested Issues

⁹ Converting examples of the MATH-type into cases of the BILL-SPLITTING-type is perhaps less intuitive but it can still be done. My very low credence that you are correct in MATH comes from the fact that all the scenarios which make you mistaken seem to me far more likely than those which make you correct. If we alter the example as to make the scenarios under which you are mistaken far less likely (perhaps by insisting that I have excellent evidence that you haven't misunderstood or misread the question) then once the disagreement comes to light I would be more inclined to think the fault lies with me. In this case the confidence gap has narrowed between us without changing my initially very high credence in the proposition.

Although there is broad agreement over how agents should respond to instances of disagreement, two related questions have divided commentators and fuelled debate. These are:

- 1) What factors affect what the rational response to a disagreement is?
- 2) To what extent should disagreeing peers revise their beliefs?

As a way of illustrating these two areas, I'll examine the direct conflict between the views of two prominent commentators, Adam Elga and Thomas Kelly.

Elga's conclusion on how to respond to peer disagreement is an account he calls the 'Equal Weight View' (Elga 2007). According to Elga's 'Equal Weight View', in cases of peer disagreement (particularly those similar to BILL-SPLITTING) those involved should give the same 'weight' to the view of their disagreeing peer as they give to their own opinion. So in the BILL-SPLITTING example I should give your belief that we each owe \$45 just the same weight as I give to my initial belief that we each owed \$43. The motivation for the Equal Weight View comes from the strong sense that epistemic peers are symmetrical in all relevant respects. The very notion of epistemic peerhood seems to require that they be so symmetrical, and so any attempt to discriminate between the views of epistemic peers will be unjustified.

Exactly how one gives a peer's view 'equal weight' is a question under some discussion, but a popular option is to 'split the difference' between the two opposing credence's. For instance, if I believe that our shares of the bill are \$43 with a credence of .1 and you believe it is \$43 with a credence of .9 then splitting the difference would mean I end up having a .5 credence in our shares being \$43 each.

Quite a lot of discussion has focused on the Equal Weight View because it is seen as an extreme view, requiring agents to radically adjust their beliefs when faced with a single conflicting peer. And in the event that the agents recognise each other as peers, the Equal Weight View requires their beliefs to conform; in other words there cannot be any reasonable and persistent disagreement among mutual epistemic peers. For this reason the charge that the Equal Weight View overstates the extent of revision required by disagreeing peers is quite common, and this is essentially Kelly's claim.

Kelly argues against the Equal Weight View by accusing it of discarding relevant evidence. The way Kelly sees it, the Equal Weight View makes responding to disagreement solely a reaction to psychological facts about what those involved believe (2005, 2010). That is, in a case like BILL-SPLITTING (even if I am in fact correct and you have calculated in error), all that I require when responding to the current disagreement is an accurate sense of my

own belief and an opinion about the reliability of your belief, so as to weigh them equally and respond. Once the disagreement has come to light, our initial calculations play no special role in determining what it is rational for us to believe, and Kelly argues this is a mistake.

Here Kelly is resisting the notion that disagreeing peers are symmetric in their disagreement. He argues instead that if one of us has performed the mental arithmetic correctly, then our disagreeing views are not symmetric, and this asymmetry should figure in rationally responding to the disagreement. So in the instance where I have calculated our bill shares correctly, this asymmetry means that I am not required to revise my initial belief to the same extent as you are.

I want to emphasise that Kelly doesn't think that when confronted with a disagreeing peer one may completely disregard their opinion. He thinks that we are probably required to cede some ground to a disagreeing peer. The main difference between Elga's view and Kelly's is that Kelly thinks that if a disputant's opinion is based on better reasons than those who they disagree with then that should figure somewhere in the rational calculus. This difference means that Elga and Kelly come to different conclusions as to the extent of revision required when one encounters a disagreeing peer. Kelly sees asymmetry where Elga sees symmetry.

Kelly's criticism says that the Equal Weight View fails to capture all the features that influence how one should rationally respond to a peer disagreement. Because of this mistake, Kelly thinks that often in cases of peer disagreement the Equal Weight View recommends ceding more ground to a disagreeing peer than is rationally required.

So, as I hope is clear, Kelly's criticisms of Elga demonstrate two of the commonly debated issues within the epistemology of disagreement and how these are connected. For Kelly, the Equal Weight View fails in capturing all features of a disagreement that are required for rationally responding to disagreement and as a result the view overstates the extent of revision required in many cases of disagreement.

Likewise for those who hold the Equal Weight View, Kelly is making a mistake in counting the quality of someone's reasons for their initial position as relevant in responding to a disagreement. Therefore, says someone like Elga, in some instances of peer disagreement agents might not revise their beliefs as much as they should.

Between stubbornness and spinelessness

When developing their accounts, theorists working in the Epistemology of Disagreement find themselves having to avoid two hazards. These hazards are the charges that their accounts make agents responding to disagreement either too eager to change their beliefs or not willing enough to change their beliefs. I'll call these two flaws *spinelessness* and *stubbornness*.

Spinelessness: Irrationally requiring/allowing disagreeing agents to change their initial beliefs.

Stubbornness: Irrationally requiring/allowing disagreeing agents to maintain their initial beliefs.

To be clear, these qualities are taken to be flaws by definition. Each of them describes a particular way in which a person can fail to meet the best standards of rationality. As such, a satisfactory account of responding to disagreement will avoid making agents either spineless or stubborn.

Spinelessness

Many of the criticisms made of analyses of disagreement can be classified as one of these flaws. In fact, Kelly's criticism of the Equal Weight View is basically an accusation of spinelessness, but concerns over spinelessness appear in other contexts as well – most notably when an account of disagreement gets applied to 'real world' examples of persistent and expert disagreement. In an example that is very close to home, philosophers worry that certain accounts of disagreement (chiefly the Equal Weight View) make their ongoing philosophical disagreements irrational.

To see this concern, take an issue on which there is deep considered debate involving many people who are expert on that issue. Typical candidates include: politics, religion, contested science and, of course, philosophy. Surely the experts working in these fields are aware that there are other experts out there who disagree with them. It also seems reasonable that the experts working in these disputed areas are peers. Recognised experts within an area are thought of as knowing as much as each other and being as intelligent as one another, at least in a broad sense.

So given these reasonable assumptions, the advice from certain accounts of disagreement (like the Equal Weight View) is that those involved in these debates should divest themselves of their strongly held beliefs and move to the middle ground. This is seen as a regrettable position to take; surely such spinelessness can't be rationally required.

Considering accounts of disagreement in the context of persistent expert disagreement illustrates the pitfalls of both spinelessness and stubbornness. On the one hand, if accounts make ongoing peer disagreement irrational, then large portions of the stakeholders in perennial debates (including many philosophers) will be irrationally stubborn. On the other hand, if we are very confident that such people can rationally maintain their beliefs, then the accounts of disagreement which paint them as irrational seem spineless. In the end, many commentators think that experts involved in persistent disagreements are not irrational and that accounts of disagreement need to be adjusted in order to accommodate them (Kelly 2007, Simpson 2012).

Stubbornness

The flip side of Spinelessness is Stubbornness. This is the concern that some accounts of disagreement don't prevent agents from irrationally maintaining their views in the face of disagreements. As such, we need restrictions built into accounts of disagreement to protect against agents acting stubbornly. A striking example of such a restriction is a principle developed by Elga and Christensen called 'Independence'¹⁰. This principle puts restrictions on the basis by which agents can assess the epistemic status of those they disagree with.

All accounts of disagreement begin with the relationship of those within the disagreement. As discussed earlier whether someone is your epistemic superior, inferior or peer is a major factor (perhaps *the only* factor) in responding to that disagreement. And so, coming to a view on someone's epistemic status is crucial in responding to the disagreement.

The threat of stubbornness presents itself in accounts where agents can freely decide what they think the epistemic status is of someone they disagree with. If we can decide the epistemic status of someone with whom we disagree without restriction, then what is to stop us from doing so in a way that is completely self-serving and circular? This is what 'Independence' is designed to prevent. In short the principle is designed to stop disagreeing agents from reasoning like this:

¹⁰ Elga doesn't use the term 'independence'. Instead he describes a procedure for rationally responding to disagreement whereby one conditionalizes on an assessment of your own reliability and your disputant's reliability which was formed prior to you thinking through the issue (Elga, 2007, p 17). In cases where priors of this type do not exist, Elga describes a process of 'factoring off' or 'setting aside'. Although they give slightly different characterizations of it, both Christensen and Elga concede that they describe a very similar notion.

'Well so-and-so disagrees with me about P. But since P is true, she is just wrong about P. So, however reliable she may generally be, I needn't take her disagreement about P as any reason at all for questioning my belief.' (Christensen 2011)

In order to prevent this type of reasoning, both Christensen and Elga¹¹ propose the following principle:

***Independence:** In evaluating the epistemic credentials of someone else's beliefs about P, in order to determine how or whether to modify one's own beliefs about P, I should do so in a way that doesn't rely on my reasoning behind my initial belief about P* (Christensen 2011).

One way to think about the principle is to compare it to basic *probabilistic independence*. If you get a temperature readout of 21 degrees from a thermometer which is historically 90% accurate, you would be justified in having a .9 credence in the temperature now being 21 degrees. If you then received another temperature readout from another thermometer, which also had a previous accuracy of 90%, and it also reported the temperature as 21 degrees then you would be justified in raising your credence in 21 degrees above .9. However, in doing so you would be making a mistake, because unbeknownst to you the second thermometer is determined by the measurement of the first. Because the two thermometers were not independent, you should not treat them as if they are independent.

The 'Independence' principle in the epistemology of disagreement is designed to work in the opposite direction. The thermometer example was a case of treating two reports as independent when they aren't, but the stubbornness worry is that disagreeing agents will treat someone's view as dependent on their own when in fact they are independent.

When I gain evidence for P, I do gain evidence that anyone believing -P is incorrect. When climate scientists observe the relationship between the increase in CO₂ in our atmosphere and a rise in global temperature they gain evidence that climate skeptics who doubt such a connection are wrong. What is the difference between this and the view Christensen is proposing?

What is true of thermometers is true of people. Say I am an avid politics junkie; I keep a close watch on the news and frequently read articles written by left leaning commentators. Because of this I have a high confidence that the Australian Labor Party will win the next

¹¹ Elga phrases his version of Independence differently, and, in keeping with the rest of his account, puts it in terms of conditionalisation (Elga, 2007). Elga says that when responding to disagreements agents cannot conditionalise on their 'detailed chain of reasoning' for their side of the disagreement.

election. If, one day, I am discussing politics with someone I have just met and they say they think Labor will lose the next election, my credence in that proposition cannot (and should not) change my estimation of their reliability about that same proposition. To do so would be to say that their reliability was dependent on my own, which seems unlikely to be the case.

Although this basic picture shows us what the 'Independence principle' is designed to achieve, some tricky questions remain about how agents stay independent across the range of cases. Most notably, accounts need to create a distinction between the aspects of a disagreement that are a threat to independence and those that should be included because they legitimately inform oneself of someone's reliability. An example of this distinction is between what Elga calls the 'circumstances of disagreement' and one's 'detailed chain of reasoning' about the proposition under dispute (Elga 2007). In the BILL-SPLITTING case, my mental calculations would be considered part of my chain of reasoning, and as such they should not influence at all my estimation of my co-diner's epistemic status. Whereas, if while doing my mental calculation, I discovered that it involves division and I know that my co-diner has trouble with division, then I should be able to use this information to alter his or her reliability.

Although this distinction seems to do good work, I am doubtful both of the distinction itself and the good work it purportedly does. I will discuss both the independence principle and probabilistic independence in a later chapter.

I will finally discuss a few of the other key notions within the larger discussion. These are: Epistemic Peerhood, Bayesianism and the Uniqueness of Disagreement within epistemology.

Epistemic Peerhood

Throughout this discussion I have been making use of the notion 'epistemic peer' without giving a full characterisation of it. So far I have relied on the platitude that epistemic peers are individuals who are neither epistemically superior nor inferior to each other within a given domain. However this is far from a substantial account, and it may not even be a genuine platitude. This is because there might be cases where agents bear no epistemic relation to one another. Such cases would be 'incommensurable disagreements' where we

are at a loss to determine anything about the disagreement, even whether or not there is a disagreement¹² because of the lack of prior/admissible information.

A survey of the disagreement literature reveals a number of different definitions of epistemic peerhood but surprisingly little direct discussion over what definition is best¹³. The lack of direct comparison between these incompatible definitions is surprising because it seems plausible that particular definitions of Epistemic Peerhood could potentially be behind some of the debate over disagreement.

As mentioned earlier, the main conflict within the disagreement debate is not *whether* epistemic peers should change their beliefs in light of peer disagreement. In fact there is broad agreement that in some circumstances those who disagree should alter their views and in other circumstances they shouldn't. What people cannot agree on is which features need to be considered when responding rationally to a disagreement. Differing views on this will in turn affect the *extent* to which those who disagree should change their beliefs in those circumstances where it seems they should.

But it seems that both of these points could be, at least to some significant degree, a matter of one's definition of epistemic peer. Say two philosophers disagree over the truth of the following statement:

'All cases of peer disagreement require giving the views involved equal weight.'

Then at least one explanation for this disagreement is that they have different definitions of epistemic peer. In fact, in the case of Elga and Kelly this is true, and by looking into these different definitions of epistemic peerhood we may discover that their disagreement is, at least partially, a merely verbal one. This is something that I will explore in greater detail in the following chapter.

Bayesianism

In many respects Bayesianism is lurking in the background of much of the recent debate over disagreement. Thinking about beliefs in terms of degrees of belief (or credences) – as well as the notion that we should change our beliefs via some process of conditioning – bear the hallmarks of Bayesianism's influence. However, although Bayesian ideas have

¹² This is in contrast with cases of merely verbal disagreement where what appears to be a disagreement over the truth of some single proposition is in reality a difference in application of terms and definitions. See Sosa (forthcoming) for some examples of this.

¹³ See Simpson (2012) and Oppy (2010) for some rare direct discussion on this question.

had a strong influence on the current debate, there have been no strictly Bayesian accounts of disagreement proposed. This is surprising because, for a Bayesian, answering the question ‘how should one respond to disagreement?’ is tantalizingly simple. As in every other case, you should conditionalise. So why aren’t those within the debate keen to embrace this simple solution¹⁴?

Taking a look at those in the debate who come close to a Bayesian account (Christensen 2009, Elga 2007, Feldman 2009, White 2009), it seems that their concern is stubbornness. In particular, Elgar and Christensen seem concerned that without the ‘Independence principle’, disagreeing agents might downgrade the epistemic status of those they should consider their peers (or even superiors) and stubbornly maintain their initial view. The implication is that something like the ‘Independence principle’ is an-add on to the basic Bayesian machinery, and without it agents could react irrationally to disagreement. But I think we should be sceptical that this add-on is needed.

In developing their accounts, these ‘near Bayesians’ in the disagreement literature utilise a very austere ‘Bayesian toolbox’; it’s a toolbox with basically only one tool – conditionalisation. However there are far more tools available to those who use Bayes’s rule alongside a host of other ideas in order to perform and analyse Bayesian Inferences. This expanded Bayesian toolset could be used to capture and make precise some of the frequently discussed scenarios and principles within the philosophical debate over disagreement.

A Bigger Bayesian Toolbox

To begin to explore how this might be done, I suggest an analogy between disagreeing individuals and the way our nervous system integrates information from different sense modalities.

In order to make our way in the world, we infer causes of all the various stimuli entering our nervous system. This information is critically important because we need to know where things are in the world around us. So when I’m in my car and I hear the sound of an ambulance, knowing which direction the ambulance is coming from will help me make a decision on how I can best get out of the way. When I hear the ambulance’s siren increasing in volume and pitch, this is my auditory system detecting the changes in the

¹⁴ There has been some work exploring some formal problems which result from a Bayesian account of disagreement (Wilson 2010, Shogenji forthcoming). Although these problems present real challenges for a Bayesian account, they aren’t necessarily fatal flaws.

siren's soundwaves and indicating to me that the ambulance is getting closer to my location.

This is an example of inferring a cause on the basis of one sensory modality (sound); however, far more accurate estimation can be made by integrating multiple sensory modalities. That is why ambulances not only have sirens but also flashing lights. If I see the distinctive blue and red flashes in my left field of vision, then that information can be integrated with the sound of the siren in order to come to a very accurate prediction about the location of the sound and the lights' common cause.

It is worth stressing just how well sensory integration works. The process of evolution is one where over long stretches of time organisms adapt in order to survive in that environment. The longer an organism has to adapt to features of their environment, the better attuned they will be to it. As a consequence, some basic features, which are advantageous in almost all environments – like the integration of auditory and visual signals – are phenomenally well tuned (Kording 2007). In fact when we compare the way human beings perform these basic tasks, we find that they combine information in a near optimal way as predicted by Bayesian statistics (Kording 2007, p. 608).

Because this process performs so well and because much of it is opaque to us, it may not immediately strike you as an instance of disagreement, but in fact there is disagreement here. Because the visual and auditory signals entering the nervous system contain a fair amount of noise, each one of these modalities creates a distribution of possible locations for its cause. It is this fact, that they offer *different* distributions, which means the final integrated inference will be so accurate. So if you adopt this way of thinking, you can see that our nervous system is in the constant business of responding to disagreements. By integrating dissimilar distributions from our various senses we can end up better off – and significantly better off at that.

The interesting thing about this process is the circumstances under which our nervous system is inclined to integrate the information from different modalities and those where it isn't. In particular, when the reports come from radically different locations, our nervous systems do not integrate these reports and infer no common cause.

If you have ever seen a ventriloquist performance, then you will be familiar with this effect. When ventriloquists perform they move a puppet's mouth in synchrony with their speech in order to create the illusion that the puppet is talking. When seeing the ventriloquist perform you are receiving reports from two different locations via your visual and auditory modalities; the visual reports come from the moving mouth of the puppet while the auditory reports come from the (comparatively stationary) mouth of the

ventriloquist. Now if these two reports are relatively close to one another then our nervous systems will infer a common cause of both the visual and auditory information and integrate the distributions from the reports. Or to put it another way, the puppet will spring to life, jabbering away. However if the two reports are far apart, say the puppet and the ventriloquist are on separate sides of the room, the illusion will disappear.

The question of interest for the epistemology of disagreement is how do we understand this difference – that two signals are integrated in one context and not integrated in another. The analogy in this case is disagreements between views, which we take seriously, and those that seem obviously wrong to us. So in the adapted BILL-SPLITTING example where my co-diner reports that we each owe \$485, quite a bit over the entire bill, I am not required to give his view anything close to equal weight.

This ‘ventriloquist effect’ has been experimentally replicated by Kording et al (2007) in order to understand how the nervous system performs these inferences and what sort of steps are needed in order to ensure their accuracy. What this work reveals is that the key step in performing an inference is deciding the appropriate model on which to interpret the reports. In the instance where the ventriloquist and puppet are close together, the appropriate model is to see the different reports as coming from a common cause, whereas when the reports come from opposite ends of the room, the reports are understood under a model where there is no common cause.

So returning to the BILL-SPLITTING case, when receiving someone’s report, a critical step is deciding what is the appropriate model under which to see my co-diner’s belief. This can be thought of as answering the question: do I really think that my co-diner and I are on the same page? Or is his belief so wild that he must be talking about something else. In light of the wildness of his view perhaps we should infer that there is ‘no common cause’ between our apparently disagreeing opinions. These questions are routinely brought up in the philosophical discussion of disagreement, but using the idea of model selection (and on from that parameter integration) allows us to give a precise characterization of what is involved in the process, for example, what sort of reports are especially relevant to model selection? Or how do one’s priors change our tendency to prefer one model over another?

Is Disagreement a Special Case?

Taking a step back and looking at the epistemology of disagreement one can quickly recognise a host of notions that are, at this stage, very familiar to modern epistemology: evidence, belief change, authority and reliability. We can, I think, expect that an answer to

the principal problem of the epistemology of disagreement will involve these notions. But given that these are already very familiar notions which epistemology has long been concerned with, why is there particular interest in disagreement?

One answer is that the interest in disagreement simply reflects pre-existing difficulties (and longstanding controversies) in understanding basic notions of evidence, belief change and reliability. Taken this way, the cases of disagreement bring relatively little to the table; instead they simply make vivid problems that occur across the range of instances involving evidence, belief change and reliability. But the way the discussion has been progressing there is little suggestion that this is the way that commentators see things.

Another explanation for the recent interest in disagreement is that the current theories of evidence, belief-change and reliability properly applied cannot adequately describe the actions of a rational agent responding to disagreement. That is, cases of disagreement offer up a *unique challenge* to existing theories of evidence and belief change and the epistemologists of disagreement's role is to develop new principles describing rational behaviour in the face of disagreements. In other words, many philosophers working in this area see disagreement as a 'special case'¹⁵.

To take an example from zoology, when echidnas and platypus were discovered in Australia, it posed an immediate problem for animal taxonomy. Prior to their discovery all known mammals gave birth to live young, and yet these strange creatures laid eggs in addition to suckling their young. For zoologists, the discovery of monotremes was a special case and it required them to act. They had to choose either to include them in the family of mammals, expanding what had previously been expected up to that point, or else create a whole new category of animal.

Special cases arise against a backdrop of our understanding of how the world works. As special cases are proposed, established schools of thought are challenged; this creates debates where the 'specialness' of the cases is a large part of what is at issue. Within philosophy we have many well-known examples. For instance, within philosophy of mind many have considered the existence of 'qualitative experience' a special case in the context

¹⁵ There is at least one noteworthy exception to this general procedure: Richard Feldman (2009) has argued that there are no novel principles of disagreement and that all cases of disagreement can be covered by a far more general principle he calls 'evidentialism'. This is the thesis that one is justified in believing P if one's evidence on balance supports P (Feldman, 2009). Feldman makes many good points concerning the state of the discussion; however 'evidentialism' as a doctrine gives little guidance as to how rational agents should respond to disagreement. It seems that while Feldman does go a significant way in arguing for the 'non-specialness' of disagreement, he does so at the expense of a detailed picture of how various pieces of evidence contribute to how one should respond.

of Physicalism. This allows contemporary dualists to claim that Physicalist accounts 'fail to capture' or 'miss out' the special case that is phenomenal experience. The Physicalist of course denies this, instead claiming that they have missed nothing, that there is no special case here and that Physicalism properly understood could account for anything the dualist wants to throw at it. This is not to claim that a Physicalist needn't concern themselves with qualitative experience. Indeed thinking about qualia has been important in helping us understand what Physicalist accounts can do and some of their potential limitations.

It is for this reason that a Bayesian account of disagreement is profitable, not only because a Bayesian account can detail how one should act in the face of disagreement, but because it also has the potential to improve our understanding of Bayesian epistemology and the ways it can be applied. It is my contention that Bayesianism, adequately elaborated, can handle anything disagreement can throw at it.

2. What is the right account of Epistemic Peerhood?

The Epistemology of Disagreement aims to give an account of the best way to respond to a disagreement. It attempts to describe what we are rationally required to do when we discover that other people think differently to us. However not all disagreements are similar, and deciding how to respond to some disagreements is easier than others.

In many disagreements the correct response to that disagreement is settled by the fact that one party possesses some epistemic advantage over the other. If I've sat down and spent a long time calculating a complex sum and you disagree with me having only guessed at the answer, it seems clear that I'm not rationally required to change my view. Likewise, because of my privileged position in the conflict, it seems reasonable that you should adjust your belief in light of our dispute.

In scenarios like this, where one disputant enjoys a clear, acknowledged epistemic advantage, it seems pretty uncontroversial how we should respond. But not all cases of disagreement are as clear. In particular, how should we respond to disagreements where neither disputant has an epistemic advantage? These disagreements, where neither party has an epistemic advantage, are problem cases for epistemologists. In order to address these problem cases, commentators have developed the notion of 'epistemic peers' – agents who are (or are taken to be) epistemically equal in some relevant sense.

Although the question of how to respond to disagreements with epistemic peers has become the focal discussion in the epistemology of disagreement there is surprisingly little direct discussion on what makes someone an epistemic peer. As Axel Gelfert points out 'comparatively little attention has been paid to a systematic analysis of what makes an *epistemic peer* in the first place' (2011, p. 508). In fact there are differing views on what makes agents epistemic peers, views with nontrivial differences and with real impact on how we think we should characterise peer disagreement.

So in what follows I will try to start to remedy the oversight by discussing two accounts of epistemic peerhood. One, which I call 'The Majority View', has been generally embraced by many commentators, and another, which I call 'The Dissenting View', has so far been less popular. In the end I will demonstrate that neither of these definitions does a perfect job of characterizing epistemic peerhood.

The Majority View

Through the course of debate a Majority View on what makes agents epistemic peers has been established (Feldman 2009, Gutting 1982, Kelly 2005, Lackey 2009, Oppy 2010, Sosa forthcoming).

For those who subscribe to the 'Majority View', A and B are epistemic peers just where they are equal with respect to their evidence and cognitive capacity¹⁶ on a given issue, at a given time. Say two doctors are diagnosing a patient. Under the majority view these doctors are epistemic peers if they have equal evidence, they have both read the same reports and test results, they have equal cognitive abilities, they are equally observant and alert, and their analytical skills brought to bear on that evidence are equal.

A few aspects of the majority view are worth clarifying. Firstly, what is meant by 'cognitive capacity' is importantly different from evidence. To get a sense of its difference, consider the goal of designing IQ tests. These tests are meant to assess one's general cognitive capacity, rather than whether or not you have some specific knowledge. They are trying to test things like memory, perceptiveness, analytic ability, intuition and so on. These capacities can relate to evidence in important ways, by either acting on evidence or influencing the evidence available to us. For example one's memory affects the amount of evidence we can retain or one's perceptiveness can effect the evidence we notice. But, as understood by the majority view, cognitive capacities are not themselves evidence for anything. They are mental tools we can use on whatever evidence we have. A bit like how a telescope tells you nothing about space until you point it at the sky, cognitive capacities become useful when applied in some way to evidence.

This separation between cognitive capacity (sometimes called epistemic virtues) and evidence allows them to stand apart from each other and vary independently. Returning to the doctor example, it seems possible that two doctors could have equal cognitive capacities whilst not possessing equal evidence. If one of the doctors had received some important test result before the other, at that moment the doctors would not be epistemic peers on the basis of their evidential disparity. This is despite the fact that they are still equal in their thoughtfulness, intelligence and diagnostic skills. Similarly, two different doctors who were basing their conclusions on the same test results but had different cognitive capacities – say if one was very tired or distracted – would not be peers even though they may have equal evidence.

¹⁶ Some commentators sometimes instead spell out the view in terms of evidence and 'epistemic virtue' (Gutting 1982, Kelly 2005). These virtues are taken to be things like intelligence, thoughtfulness and observance etc. These are the same abilities which I'm including under the banner of 'cognitive capacities'. Call them 'epistemic virtues' or 'cognitive capacities', it makes little difference.

Another aspect of the majority view is that epistemic peerhood is situation specific, meaning it is sensitive to particular topics and even perhaps particular propositions. Two agents may be epistemic peers on a particular topic but not epistemic peers on another topic. Just because you are my epistemic peer on questions of basic arithmetic you need not be my peer in identifying celebrities. Similarly agents might be peers at a particular time but not at another time. It isn't enough that agents have equal evidence and cognitive capacity in general; they must have similar evidence and cognitive capacity with respect to the specific issue at hand.

The majority view has a lot of intuitive appeal because 'evidence' and 'cognitive capacity' do seem consistently epistemically relevant. Whatever the subject, our beliefs about that subject will be the product of our evidence and the cognitive skills we bring to that evidence. If these are the relevant factors in producing our beliefs, it makes sense to say that if individuals are alike in these respects then neither has an epistemic advantage over the other.

Ambiguity in The Majority view

Even though the parameters for evidence and cognitive capacity seem like good criteria for epistemic competence, there are difficulties in saying precisely what it means to have equal evidence and cognitive capacities.

Like most philosophers working in the epistemology of disagreement, those who hold the majority view are concerned with what makes two agents epistemic peers because it sheds light on how agents should respond to disagreement generally. However among those who hold the majority view, there is ambiguity on exactly what role the *fact* of epistemic peerhood plays in shaping how one should respond to disagreements among epistemic peers.

It is possible, under the majority view, that I disagree with someone who is in fact my epistemic peer by virtue of them being my evidential and cognitive equal, and yet I will not know they are my epistemic peer. Say I'm at a dinner party with a bunch of people I've only met that evening and we are discussing film trivia. Eventually a disagreement is uncovered between myself and one of the dinner guests, someone I've only just met that evening. Now I know myself to be a real movie buff. I'm correct in believing I know much more about movies than your average person, and so I believe that the particular person who disagrees with me now is not my epistemic peer. But as things turn out, I'm wrong. The person who disagrees with me is in fact my epistemic peer: they know just as much as

I do about movies and their cognitive capacities are in fact equal to my own. Although I don't believe they are my epistemic peer and I have no particular evidence to suggest they are, the fact that they have equal evidence and cognitive capacities makes them my peer.

Does the fact that this person is my epistemic peer make my dismissal of their views unjustified? If it does, it's difficult to see how. At the time of our disagreement I have far more evidence that they aren't my epistemic peer; knowing as much about movies as I do is very rare and in my personal history I've never met anyone who knows as much (or more) about film as I do. So given my evidence up to this point, the view best supported by my current evidence is that someone I meet at a dinner party is my 'movie-knowledge inferior'. Being unjustified by simply responding to my current evidence and being required to react to epistemic status when that status is unsupported by my current evidence is a hard position to hold.

It seems that when responding to disagreement what matters is not the fact that two people are epistemic peers but the extent to which they believe they are epistemic peers¹⁷. This is a feature that is rarely acknowledged by those who hold the majority view. Once more, if this is the case, the definition of epistemic peerhood advocated by majority view holders is more or less irrelevant to how one should respond to cases of peer disagreement, that is, unless a case can be made for how the mere fact of epistemic peerhood can and should influence our beliefs about who is an epistemic peer.

Perhaps such an externalist account¹⁸ is possible; in order to progress, let's suppose that it is. In this case then we will return to the issue of saying exactly what sort of equality makes for the fact of epistemic peerhood. At first blush the majority view has an appealingly simple answer to this; epistemic peerhood is determined by equality of 'cognitive capacity' and 'evidence'. But by drilling down into these categories it's less than perfectly clear what these categories refer to. As such what makes for equality in these features is still an open question.

Cognitive Capacity

¹⁷ Now it may turn out that this is the position of many of the majority view holders, that what they mean when they say that epistemic peers are equal in evidence and cognitive capacity is just that epistemic peers are those people who are thought to have equal evidence and cognitive capacity. But if this is the case, I can see little advantage in making the distinctions the majority view makes. It seems similarly vacuous to say that an epistemic peer is someone who you think is your epistemic equal.

¹⁸ By 'externalist account' I mean an account whereby the mere fact that agents are cognitive and evidential equals impacts how they should respond to a disagreement. This contrasts with a more subjectivist account where the belief about the epistemic status of those involved in the disagreement would determine how they should respond.

One way to compare individuals' cognitive capacity is to start by breaking down cognitive capacity into more basic capacities and comparing these between agents. Understood this way, saying that what it means for agents to be cognitive equals is that they must be equal in their intelligence, reasoning skills, memory recall and so forth (Kelly 2005 pp. 174–175, Oppy 2010). But this only serves to push the burden of analysis onto these more basic cognitive capacities. This move leaves unanswered what makes two people have equal intelligence, reasoning skills or memory recall etc. It is true that some of these aspects of cognitive capacity are easier to characterise than others. It is far easier to give a precise definition of a specific capacity such as 'memory recall' than a very complex combination like 'intelligence'. But perhaps the most complex cognitive virtues, like 'intelligence', can be sufficiently reduced so that all of their component parts can also be given a specific, testable definition. This would in turn make comparing agents for cognitive equality a testable enterprise. If two agents perform equally well at a battery tests designed to measure these basic 'atomic' cognitive abilities then they are cognitive equals.

However this raises the further question of what makes a particular psychological feature (basic or complex) included in the categories of 'cognitive capacities/epistemic virtues'. I think the best way to understand cognitive capacities is as non-evidential features of someone's psychology that affect one's chance of forming accurate beliefs. Intelligence is regarded as an important epistemic virtue because, all other things being equal, people who are more intelligent are more likely to have accurate beliefs than less intelligent people. The reason 'rashness' is thought of as an epistemic vice and 'methodicalness' is an epistemic virtue is because methodicalness increases one's chances of having accurate beliefs, whereas rashness decreases those chances.

If raising or lowering one's propensity to form correct beliefs is what makes a psychological feature an epistemic virtue or vice, then why bother defining epistemic peerhood in terms of cognitive capacities at all? It would perhaps be more direct to instead deal with the raising and lowering of one's chances to form true beliefs, instead of dealing with an estimation of these chances in the form of cognitive capacities.

I think the reason why cognitive capacities are so attractive is because they are generalised capacities. Something like intelligence is a general capacity: when someone displays intelligence in one domain, we assume that their intelligence will generalise to many other scenarios. This feature of cognitive capacities is important in coming to *judgments* of someone's general reliability. The fact that we can observe someone acting intelligently, thoughtfully or methodically in one situation gives us a good indication that they will act

similarly in some other, quite different, situation¹⁹. But there is a risk in thinking that a cognitive capacity applies to a situation where it in fact doesn't.

How big a risk we are taking when we characterise epistemic peerhood in terms of cognitive capacities depends on how often we correctly generalise those capacities. It is possible that the insistence on general cognitive capacities encourages us to generalise more than we would otherwise and therefore increase the risk that we might generalise incorrectly.

On the other hand perhaps there is no harm in characterising an agent's chances of forming correct beliefs through their cognitive capacities. If cognitive capacities are a way of organising and cataloging someone's propensity to have accurate beliefs, then there may be no harm in using them to characterise epistemic peerhood. If the link between one's propensity to have accurate beliefs and one's epistemic virtues is direct, then there may be no particular harm, but there will also likely be no benefit and it would be unnecessary. If what makes a particular property an epistemic virtue is its tendency to yield correct beliefs in the owner, then the majority view holder needs some account of why it is beneficial to characterise epistemic peerhood in terms of these virtues.

Evidence

Although the majority view says epistemic peers are those who have equal evidence, similarly to cognitive capacity, rarely is any detail given on what determines evidential parity.

One way would be to say that individuals are evidentially equal when they have the *same evidence*. Sameness of evidence would be achieved by gaining the same reports via the same channels. For example, two cognitive peers watching election result coverage on television have the same evidence for who has won the election because they are receiving the same reports via the same reporter.

Given this account there will be boundary cases which will test the notion of 'sameness of evidence'. Suppose two doctors look at separate copies of a patient's blood test results. Provided both copies report the same results for the test, the doctors have the same evidence. Reading the same result from a different printout is a trivial difference, one that doesn't make for difference of evidence between the doctors. However if it is possible that one of the printers made an error when printing the test results, the reports the doctors

¹⁹ This is sharply in contrast to evidence, which is highly non-generalizable.

receive will be different. In this circumstance the doctors no longer have the same evidence, but can the different evidence still be regarded as equal?

It isn't clear from those who hold the majority view whether these doctors would be evidentially equal or not.

On the one hand the doctors are clearly evidentially different from one another. One of them has received a reliable report while the other received a misleading one. The question is; does this difference in evidence necessarily make them evidentially unequal. On the face of it we would not want to say that an individual with misleading evidence is equal to someone whose evidence was not misleading. We would instead say that people with misleading evidence are evidentially inferior to those whose evidence isn't misleading.

On the other hand it is hard to see how the mere fact of a misleading report is meant to influence someone if they have no particular evidence for the report's misleading nature. Say we deem the unlucky doctor as evidentially inferior and then the two doctors discuss this patient's test results. Even if the doctors are cognitive peers the unlucky physician will be epistemically inferior to her colleague because of her lower quality evidence. This is despite the fact that neither doctor has any evidence that they are not epistemic peers. So when the disagreement surfaces it seems strange that one doctor's opinion should hold sway over the other, given the lack of any information suggesting that they are not epistemic peers. Recall that we are developing an account of evidential parity in order to say how we should respond to cases of disagreement among epistemic peers. In the case of a doctor who receives a faulty printout, if she doesn't have any additional evidence that the printout is faulty, how would this external fact affect what is rationally required of her?

This dilemma follows closely the general tension within the majority view: to what extent should epistemic peerhood be a function of external factors the agent may be unaware of, or to what extent should epistemic peerhood be based on features of an agent's belief set. We will return to this question in following sections.

However, to best answer this tension, one thing, I think, is relatively clear: characterising equality of evidence as 'sameness of evidence' is far too restrictive. Because while sameness of evidence is sufficient for equality of evidence, it should not be necessary; there are many situations where individuals may be construed as evidential equals despite having different evidence.

Consider a toy example where you and I are each given half the pieces of a jigsaw puzzle. It's a special puzzle where the image is equally spread throughout its many pieces so that each piece carries an equal share of what the complete puzzle image is. We aren't allowed

to see each other's pieces so our pieces comprise a unique set of evidence for the image depicted on the completed puzzle, and we have half of these each. In this scenario I think it makes sense to say that although we possess entirely different information sets by virtue of having access to entirely different sets of puzzle pieces, we do have equal evidence for our beliefs in what the completed puzzle depicts. This is because our evidence (our sets of puzzle pieces) lends equal support to our beliefs in what the complete puzzle depicts.

If I was to have 90% of the puzzle pieces and you were to have only 10%, then clearly I would be evidentially superior to you. But given that we have an equal number of puzzle pieces, our evidence, although different, still has the relevant parity.

This observation gives some clue to an account of evidence that can be used in the majority view – that what is important when comparing agents for evidential parity is that we are not comparing the evidence itself but rather we are comparing the role of evidence. More specifically we are comparing the extent to which a given piece of evidence confirms or disconfirms a particular proposition. Agents whose evidence equally confirms/ disconfirms can be thought of as evidential equals despite the fact that they may be aware of different propositions.

Elga's alternative view

Adam Elga (2007, p.13) advocates an alternative dissenting view on epistemic peerhood to the majority view. Elga argues that A counts B as an epistemic peer on an about-to-be-judged claim if, conditional on their disagreeing about P, A thinks that B is equally likely to be mistaken about P.

So, for example, say you and I are out bird watching looking for a rare forest bird. According to Elga, you are my epistemic peer in identifying this bird if, in the event that we disagree on the species of a particular bird, I think that you are as likely as I am to be mistaken about that bird's species.

Elga offers the following justification for favouring his account over the majority view: imagine that you consider someone less likely to be correct on a given issue, then according to Elga, it would be odd for you to regard them as an epistemic peer regardless of their evidence or cognitive capacities (Elga 2007). If while bird watching, I consider you to be far more likely than myself to be mistaken on a given issue then, regardless of your evidential or cognitive capacities, it would be strange for me to think you are my epistemic peer.

This observation is not likely to convince a majority view holder that they are mistaken. At best, they might take it as evidence that one account entails the other. Even if a majority view holder accepts Elga's justification, they could respond by arguing that the only reason you would think someone is equally likely to be mistaken as yourself – given a disagreement – is because they *must be* your evidential and cognitive equal. That is to say that it is impossible for you to think someone is just as likely to be mistaken about something whilst also thinking they aren't your evidential or cognitive equal. This view falls out of a background notion which I think is common to the majority view holder. This is the notion that evidence and cognitive capacity are the sole contributing factors for agents justifiably believing as they do and also the only factors that contribute to someone's propensity for accurate beliefs.

It is a claim that amounts to saying that the majority view entails Elga's definition of epistemic peer – that someone's chances of having correct beliefs are determined entirely by their evidence and their cognitive capacity. If I am someone's evidential and cognitive equal, I must consider him or her equally likely to be mistaken on a disputed issue. However I think we have good reasons to doubt this claim.

It is possible that two agents are evidential and cognitive equals without counting as epistemic peers under Elga's definition. This is because, as I've mentioned earlier, the majority view is concerned with epistemic peerhood as determined by objective features of individuals. Evidence and cognitive capacity are the sorts of things which agents possess regardless of their own or anyone else's beliefs about them; as such an account of epistemic peerhood that is based on these notions will be an objective matter not dependent on anyone's beliefs. This is clearly not the case for Elga's account, which makes epistemic peerhood wholly a question of someone's beliefs. So it isn't the case that the majority view entails Elga's account of epistemic peerhood.

But perhaps instead the entailment goes the other way – that Elga's account entails the majority view. Again I think there are good reasons to say it doesn't. It seems perfectly possible that someone be judged an epistemic peer under Elga's definition without meeting the standards of evidential or cognitive equality. As I have argued earlier in this chapter, a basic notion of the majority view is that evidence and cognitive capacity can vary independently so as to make someone epistemically better (or worse) off. If two people are equal with respect to their cognitive capacity but one has more evidence, then that person will be epistemically better off. Similarly, if agents are equal in their evidence but one agent has superior cognitive abilities, the agent who is better able to analyse and correctly understand the evidence will be better off.

Presumably we could tweak these parameters so that agents could differ both in their cognitive abilities and their evidence and yet they could still be equally likely to be mistaken on a given topic.

Say you and I are participating in an experiment where we are trying to memorise flash cards in a given time. At first we are each given the same amount of time to memorise 10 flashcards. It turns out that you can memorise an average of 8 flash cards in the given time whereas I can only memorise 4 on average. From this initial experiment we could infer that given the same amount of time (or evidence), you have a superior working memory to me and are therefore my cognitive superior. The experimenter then compensates for your superior visual memory by adjusting the time we each get to see the cards. Your time is progressively decreased and mine is increased until on average we are each memorising 6 cards at a time. In this scenario it seems that neither of us have equal evidence or equal cognitive abilities, but because these parameters are so arranged we still may be epistemic peers under Elga's definition.

This scenario may seem artificial but I can see no reason why evidence and cognitive ability couldn't be similarly balanced in real world disputes so as to make disputants equally likely to be mistaken on a given topic.

But perhaps this is to misunderstand the majority view. Perhaps it is not important that epistemic peers have equal evidence and equal cognitive ability but just that these factors combine to produce agents that are equally well placed on a given topic.

Whereas previously we had taken the majority view to be expressing the following:

Two agents, S1 and S2, are epistemic peers iff:

$$1) S1(\text{evidence}) = S2(\text{evidence})$$

AND

$$2) S1(\text{cognitive capacity}) = S2(\text{cognitive capacity})$$

Perhaps we should read the majority view as saying instead:

Two agents, S1 and S2, are epistemic peer iff:

$$1) S1(\text{evidence and cognitive capacity}) = S2(\text{evidence and cognitive capacity})$$

But this seems to undermine part of the initial motivation for the majority view that these particular parameters play a special role in determining epistemic status. If epistemic peerhood is not measured by parity in evidence and cognition but instead in how these

combine (perhaps with yet more parameters) to produce a general epistemic score, perhaps this should determine epistemic status. If the majority view is going to be a contender, its defenders need to give some justification for the special status afforded evidence and cognitive capacity.

A striking difference between the majority view and Elga's account is that Elga's account makes it possible that the epistemic peerhood relation be asymmetric. That is, Elga's account makes it possible for Alfred to count Betty as his epistemic peer even though Betty does not count Alfred as her epistemic peer. In the majority view, although Alfred and Betty can have differing attitudes toward their epistemic relationship, there will still be a fact of the matter as to whether or not they are epistemic peers, and this fact is unaffected by their beliefs about their epistemic relationship.

Not so in Elga's account. According to Elga, counting someone as an epistemic peer is all there is to being an epistemic peer. Since it is clearly possible for two agents to have different views on their status as epistemic peers, it is equally possible for epistemic peerhood to be asymmetric.

Problems for Elga's account.

There are a few purported problems for Elga's account of epistemic peerhood. To begin with, Elga's account seems far less applicable to the range of cases (Oppy, 2010, p. 188). As such the account restricts the application of epistemic peerhood to a far narrower set of instances than those we might want to consider as epistemic peers. If an account of epistemic peerhood fails to assign the relation in situations where we intuitively think it should, then so much the worse for that account.

In particular, because Elga's account concerns 'about to be judged claims', his account says nothing in cases where agents have already come to a view on a particular topic. To use one of Elga's examples, say you and I are at the horse races and we are about to watch the first race of the day. At this moment, before the race begins, I consider you just as good at judging horse race winners as I am. I think that if we were to disagree we would be equally likely to be right (or wrong) about who won that race and therefore, under Elga's account, we would be epistemic peers.

We would be epistemic peers, that is, up until the end of that first race. Once the race has been run and we have each come to a conclusion over which horse won it, then we would no longer be epistemic peers. Elga's account specifies that epistemic peerhood is an attitude toward an 'about-to-be-judged' claim, and since the claim in question no longer

falls into the 'about-to-be-judged' category, Elga's account has nothing to say on our peerhood status. A similar problem arises in cases where the disagreeing parties formed their divergent judgments years earlier, long before they ever met each other. In this case those involved may fail to have a prior view on the reliability of the other party at the time the judgment was formed.

On the face of it, this does seem like a deficiency in Elga's account. The need for an account of epistemic peerhood arose because we needed some standard for what makes people equally good judges on some issue. Why then should the notion not apply to cases that have already been judged? The restriction does look puzzling.

However, a part of Elga's account does attempt to address this problem. Elga's defining of epistemic peerhood asks you to take what you know of your own epistemic resources and compare them with what you know of someone else's epistemic resources. If you consider these equivalent and you think that you each stand an equal chance of being correct, then you are epistemic peers. In situations where you don't have any pre-disagreement views on a disagreeing party's epistemic credentials, then Elga suggests you compare your current epistemic position with your disputant's position, but you must do so 'setting aside' your reasoning for your own opinion (Elga, 2007, p17).

Elga gives very little detail of how precisely one 'sets aside' their own reasoning in the face of a disagreement. And although the idea seems intuitive enough, to establish whether or not it unproblematically allows Elga to extend his account of epistemic peerhood requires carefully working through all details. This is something that I will do in my next chapter.

Another possible problem for Elga's account is to do with its theoretical underpinnings. Elga's account of epistemic peerhood is strongly influenced by Subjective Bayesianism – seeing epistemic peerhood as a particular subjective conditional probability judgment. More precisely, it is the evaluation and comparison of two subjective probabilities: the probability that I am correct conditional on a disagreement and the probability that you are correct conditional on that same disagreement. If, for me, these two conditional probability values are the same, then I count you as an epistemic peer.

However there is perhaps a problem for Elga's account if we take to heart the consequences of Subjective Bayesianism; in particular, Bayesianism's interpretation of probability complicates Elga's account.

Subjective Bayesian's hold that probabilities are degrees of belief. So when I say that the chance of rolling an even number on a normal die is $\frac{1}{2}$, I am not making a statement about what would occur over a series of similar dice rolls (either actual or hypothetical); I am

instead revealing that the extent of my belief in rolling an even number is .5. What I think of as the chance of a particular event is my degree of belief in that event occurring, otherwise known as my *credence*.

Part of what makes Subjective Bayesianism an appealing notion is its ability to generalise from being an account of chance to being a general epistemology. Modelling beliefs on probability axioms give epistemologists powerful resources to solve epistemological problems (Hájek, & Hartmann, 2010, ppg. 101-102).

The key complication for Elga's account when dealing with credences is: what makes someone's credence 'equally likely to be wrong' as someone else's credence? If I have a .8 credence in P and you have a .7 credence in P, we clearly have a disagreement: we have different confidence in P being true. But how could I think that you are just as likely as me to be incorrect about P?

If you were an objectivist about chance, then the answer would be relatively simple. Because the chance of P would be an objective fact about the world, at least one of us would be wrong, and if I thought that we were equally likely to be correct over the objective chance of P, then we would be epistemic peers.

However for the subjectivist about chance, things are less clear. Returning to the racetrack again, say you and I have both studied the racing form guide and as a result we come to different credences for Horse-A winning; you think it has a .7 chance of winning the first race, and I think it has a .8 chance of winning the first race. If we then reveal these credences to one another, how do I come to a conclusion on how likely you are to be correct about this?

Say we each placed bets on the horse winning in accordance with our credences: you place a \$8 bet and I place a \$7 bet. In the event that the horse loses, then perhaps I have reason to think that your belief is worse than mine because you lost more money. Looking at things this way at least gives us some scope to compare credential beliefs, but the comparison being made isn't similar to Elga's initial account. We have stopped comparing our beliefs for their propensity for accuracy and instead we are now comparing them for pragmatic cost.

Also when looked at this way, we find that the only circumstance where one agent could think that someone else's credence stands an equal chance of being correct as their own is if they think that their credences are the same (Oppy 2010). Perhaps the only circumstances where credential beliefs have an equal chance of being wrong are when they

are in fact the same! Seeing this there is a genuine concern that Elga's account makes disagreement among epistemic peers impossible.

Conclusion

The epistemology of disagreement is concerned with two questions. Firstly, what makes two agents epistemic peers, and secondly, how should a rational agent respond to a disagreement with an epistemic peer?

The majority view offers an appealing answer to what makes two people epistemic peers. It presents an objective account of epistemic peerhood that allows for any two people to be compared for peerhood. Even if two people have never met, the majority view provides grounds for determining whether or not they are in fact epistemic peers.

Although this view has its benefits, it is unclear how the fact of a peerhood relationship is meant to influence how a rational agent should respond to a given disagreement. Unless the majority view holder can offer an account of how their conception of epistemic peerhood influences how one should rationally respond to disagreement, then their account does almost no real work.

Elga's dissenting account, on the other hand, is more geared toward coming to an answer on how to rationally respond to disagreement. Casting epistemic peerhood as a subjective judgment of someone else's reliability under certain conditions makes it far easier for the notion to play a direct role in shaping how a disagreeing party should respond.

However Elga's account has its own problems. In particular Elga's account seems to be overly restrictive when it comes to applying the notion of epistemic peerhood. But according to Elga, these restrictions are there for a good reason; they prevent agents from unjustifiably reassessing the epistemic status of those they disagree with. But is Elga right to worry about this type of thing? Are the restrictions he (and others) make on how rational agents should respond to disagreements necessary? I don't think they are.

This point is larger than the issue of defining peer hood and it goes more directly to a question in my introduction: is disagreement – even peer disagreement – really a special case? If it turns out that the something like Elga's restrictions are justified then we might have reason to think that disagreement is a special case, but if they can't be justified why not remove them and endorse a conception of epistemic status which as general and as widely applicable as possible.

In what follows I will examine more closely two reasons why disagreements might be considered 'special cases', including discussing some of the unique inference rules discussed by Elga.

3. Is Disagreement special?

Higher Order Evidence and inference rules

Introduction

As mentioned in the preceding chapter, the recent debate over disagreement has proceeded on the basis that disagreement is something of a 'special case'. By special case I mean that the cases of disagreement most often discussed cannot be easily accommodated by our general epistemic theories, making disagreement a problem case for those theories.

What are some of the reasons why disagreement might be special? Surveying the literature there are two potential reasons why disagreement might be considered a special case. One is that disagreement involves a special type of evidence, namely Higher Order Evidence, and the second is that disagreement requires special inference rules. In what follows I will discuss each of these reasons and I will argue that they do not constitute a convincing case for the specialness of disagreement within contemporary epistemology.

Disagreement Involves Special Evidence

The first claim I'll explore is that cases of disagreements distinguish themselves because they involve a special type of evidence, namely Higher Order Evidence.

Precisely how to characterise Higher Order Evidence is not uncontroversial. But to introduce the idea, consider the following example borrowed from Christensen (2010):

A first year medical intern is working at a hospital and has just finished a shift, during which it was their task to admit new patients, observe their symptoms and make diagnoses based on that evidence. The intern is leaving an interview room having just diagnosed a new patient when a friendly nurse approaches and informs them that they've been awake for 36 hours. Appreciating how being sleep deprived can impair someone's cognitive skills they now become far less confident in the diagnosis they had just made and resolve to make sure another doctor double checks their patients.

In this story the steps taken by our first year intern seem reasonable and prudent. If you were the patient just seen by the intern you too might want a second opinion concerning your diagnosis. But it's not immediately clear why this is the case. What does your doctor's lack of sleep have to do with your illness? Them being awake so long doesn't change your symptoms; it doesn't make your coughs worse or your white blood cell count

higher. So what is the difference between a patient's symptoms, which clearly have a direct impact on diagnosis, and the news that the young doctor is dramatically sleep deprived?

The difference can be understood as the difference between First Order and Higher Order Evidence. First Order Evidence is what we often mean when discussing evidence: it is that information which makes the truth of a proposition more or less likely. So in our example, in relation to a particular diagnosis, the patient's symptoms are First Order Evidence because they impact directly on which illness the patient has. Contrastingly, the news that the doctor is sleep deprived is Higher Order Evidence because this information doesn't directly affect the chance of a patient having a particular illness; instead it influences the chance that the doctor has correctly understood the first order evidence available. The first order evidence itself isn't in doubt; what is in doubt is the doctor's appreciation of that First Order Evidence. So speaking broadly, Higher Order Evidence is very often that evidence which bears on someone's ability to assess/interpret evidence correctly²⁰.

In cases of disagreement, Higher Order Evidence features prominently in two ways. The first way is in the judgment agents make about one another's epistemic status/credentials. When anyone discovers that they disagree with someone else, how they respond to the disagreement depends crucially on comparing the judged epistemic status of the disagreeing parties. Disagreeing with someone who I regard as an expert will prompt a very different response than disagreeing with someone who I think is a charlatan. These judgements about someone's epistemic credentials are themselves based on evidence. Relevant evidence might commonly include things like their track record for accuracy and honesty or their endorsement by other experts. In all cases the evidence that bears on someone's epistemic status indicates how seriously one should take their opinion.

Exactly which evidence is permissible when forming a judgement about someone's epistemic status is an interesting question and one that I'll discuss in the second part of this chapter. For now it's important to recognise that the evidence that bears on someone's epistemic status is Higher Order Evidence because it is evidence about evidence. Specifically, it is Higher Order Evidence that shows us how to weigh the first order evidence present in the opinions of others. When someone avows that they believe P, then

²⁰ This is by no means an exhaustive definition of Higher Order Evidence. Richard Feldman and Thomas Kelly use a far more permissive definition where Higher Order Evidence is 'evidence about the existence, merits or significance of some evidence' (Feldman, 2009, 304). Christensen however takes a much narrower view, very similar to what I've said here. And seeing as it is his views that will preoccupy us for much of this chapter, I will adopt his usage.

whatever information indicates that they are in fact correct about P is Higher Order Evidence with respect to P.

The second way Higher Order Evidence features prominently in disagreements are in the disagreements themselves. Disagreements between epistemic peers are particularly striking cases of disagreements presenting Higher Order Evidence to those involved. According to the most popular conception, disagreeing with an epistemic peer means disagreeing with someone who you think of as having the same evidence and having equal ability in assessing that evidence as you do (Feldman 2009, Gutting 1982, Kelly 2005, Lackey 2009, Oppy 2010, Sosa forthcoming). Say you and your peer were faced with a simple arithmetic problem, you can both see the problem clearly, have both had the same amount of thinking time and you have equally good track records in solving problems like this. If you then discovered that your peer disagrees with you about the solution, it seems very intuitive that this discovery should increase your confidence that you may have made a calculating mistake and consequently you should lower your confidence in your solution. Like in the medical intern case, disagreeing with your peer lowers your confidence in your answer not by presenting you with first order evidence about the arithmetic but by presenting you with evidence that you have mistakenly assessed some evidence you already possess. The fact that your peer disagrees with you is Higher Order Evidence because it suggests that you failed to correctly interpret your shared evidence.

The fact that Higher Order Evidence plays such a prominent role in disagreements gives rise to a quick argument for the specialness of Disagreement within the epistemic landscape. I imagine it would go something like this;

- 1) Higher order evidence is a special evidence.
- 2) A set of cases is epistemically special if they involve special evidence.
- 3) Cases of disagreement involve large amounts of Higher Order Evidence.

Therefore:

- C) Cases of disagreement are special.

Higher Order Evidence is special because it is importantly different from first order evidence. Also, Higher Order Evidence features prominently in cases of Disagreement. Therefore cases of disagreements are special at least to the extent that they feature large amounts of Higher Order Evidence. In what follows I will try and show that this quick argument is false by undermining the justification for the first premise that Higher Order Evidence is importantly different from first order evidence. To begin with we need to

explore some of the proposed reasons why Higher Order Evidence might be importantly different from garden variety first order evidence.

Why might Higher Order Evidence be interestingly different from First Order Evidence?

In his 2010 paper 'Higher Order Evidence', David Christensen raises a number of reasons why Higher Order Evidence is interestingly different from the First Order variety. In what follows I will discuss each of these reasons and find that each fails to make a convincing case for the distinctiveness of Higher Order Evidence.

Candidate 1: Responding to Higher Order Evidence requires cognitive fallibility.

As believers who want accurate beliefs, we can fail in two ways. One way we can fail is by forming our beliefs on the basis of limited evidence. We are finite observers and we gather evidence in a piecemeal way – gaining more evidence, piece by piece, as time goes by. At every occasion, as we gather new evidence we abandon our previous position on whatever the evidence is about. This process acts as a constant reminder that we form our beliefs in the context of incomplete evidence. The gaining of new evidence throws into sharp relief just how limited our previous set of evidence was compared to the evidence we now possess.

This perhaps sounds regrettable, but we probably shouldn't be too hard on ourselves about it; after all, in order to be rational one doesn't need *all* the evidence. For many, rationality instead comes down to merely respecting and properly understanding the evidence you do have. Rationality doesn't require us to be omniscient; it just requires us to do our due diligence according to whatever evidence we possess at any particular time.

The other way we can fail to have accurate beliefs is by misunderstanding the evidence we do have. Compared to merely having incomplete evidence, this failing seems more blameworthy. In contrast to just having limited evidence, the idea that we misinterpret the evidence we have already gathered stands in direct threat to our rationality. If rationality requires us to respond to evidence correctly as it is gathered, then failing to do this (i.e. to respond inappropriately to evidence, misunderstanding it) is to fail in the most basic requirement of inductive rationality.

To see this point more clearly Christensen (2010) asks us to consider an agent that is 'cognitively perfect' (p. 191). Such an agent has perfect deductive and inductive insight, always drawing the appropriate credence's from whatever evidence she has at the time. If a cognitively perfect agent has some evidence, and in response confidently believes P, her

credence is the optimum response to the evidence she currently has. Given her evidential circumstances she could not draw a better conclusion. She then gets strong Higher Order Evidence that suggests that she made an error in her previous evaluation of evidence. This higher order evidence would, of course, be misleading evidence as it is impossible for her to make such an error, but a cognitively perfect agent is not immune from misleading evidence. As Christensen says, misleading evidence is precisely that evidence which rationally leads one away from the truth (Christensen, 2010, p. 191).

Having encountered misleading Higher Order Evidence, a cognitively perfect agent would assess the importance of misleading evidence perfectly and become far less confident in P than they were originally. Christensen argues that when responding to first order evidence the cognitively perfect agent is essentially blameless. Correctly responding to the evidence gathered so far is really all that we can ask from an agent trying to form rational beliefs. The alternative is that they form their beliefs against or outside the evidence that they currently possess, an idea that is antithetical to reasonable belief formation.

And yet, Christensen argues, responding to Higher Order Evidence forces even a cognitively perfect agent to 'embody a kind of epistemic imperfection' (Christensen, 2010, p.193). This is because even though the cognitively perfect agent flawlessly grasps the import of her first order evidence, the encountered higher order evidence prevents her from maintaining beliefs that utilise her perfect evidence-assessing capabilities. Instead she is compelled to have a first order belief (i.e. a belief about P) not based on the best possible interpretation of her first order evidence (i.e. evidence E). And that does seem like a waste of a perfect evidence assessor.

The important aspect of this discussion is that it is supposed to demonstrate the distinctiveness of Higher Order Evidence from First Order Evidence. The idea is that responding to Higher Order Evidence forces us into epistemic imperfection while responding to First Order Evidence does not. But is this the case? Does Christensen's discussion really show that responding to Higher Order Evidence involves a failing that is distinctive from the failing of limited evidence? In what follows I will argue that it doesn't and that the epistemic failing raised by Higher Order Evidence is the same as was present in reacting to First Order Evidence.

Candidate 1 response: Responding to Higher Order Evidence doesn't require cognitive fallibility

Recall the idea that when responding to first order evidence, it is possible to remain epistemically blameless, that although we change our beliefs in response to new evidence

we can still believe whatever our total evidence warrants at any particular point in time. Even though we do not currently believe what we once believed, we don't look back and necessarily see a history of errors in evidential evaluation. It is as though with every new piece of evidence we may rightly say to ourselves: 'According to what I know now things aren't as I thought they were, but according to what I knew then I shouldn't have believed anything different'. My argument is that this is essentially the same in cases of responding to Higher Order Evidence – that in responding to both First Order and Higher Order Evidence the only epistemic limitation we admit to is that we form our beliefs on the basis of limited evidence, and as discussed earlier this is not a great epistemic failing.

The first point to note is that all of the examples of Higher Order Evidence discussed by Christensen happen to be ones where the encountered Higher Order Evidence bears on some *previous* judgment of evidence. The medical intern case is an example of this; the information that the intern is sleep deprived affects his confidence in previous diagnoses. However there is nothing in the basic definition of Higher Order Evidence that requires it to be 'back-facing' in this way. Indeed much of the Higher Order Evidence one receives will not only bear on one's previous abilities in interpreting evidence but also one's present and future abilities. If after having learnt about their sleep deprived state the medical intern makes another diagnosis, their confidence should be lower – given the just encountered Higher Order Evidence – than it would have been if they had never been informed of their lack of sleep. Furthermore, when Higher Order Evidence bears on future evidence, properly responding to that evidence doesn't require any particular rational failing.

Christensen focuses his discussion of Higher Order Evidence on that information which informs us about our own cognitive capacities. If we appropriately understand this evidence it will have a significant effect in altering the beliefs we *then go on to form*. For instance if I gain evidence that I am not as cognitively capable as I normally am, say by realising I'm drunk, then I will rightly be less confident of all the conclusions I draw thereafter.

Say I am at a dinner party and I realise I'm a bit drunk. It seems at least possible that I could realise my drunken state and properly account for it by lowering my confidence in the conclusions I come to *as I form them*. And it seems possible that I could do this so well that all the beliefs I formed whilst being drunk are the same as those I would have come to if I'd only found out about my drunkenness the next morning and then had to revise beliefs which I now think are overconfident. My point is that Higher Order Evidence can (and almost always does) contribute to our first order beliefs as they are formed. And in

these more mundane situations we don't think that responding to Higher Order Evidence commits us to any sort of epistemic failing.

The truth is that whenever we form beliefs we do so with some understanding of our ability to correctly interpret evidence. For example, the hotshot math student is very confident that they have correctly understood an exercise question, applied the correct tools and gotten the right answer. This is because they have a wealth of Higher Order Evidence supporting their confidence; they are accustomed to getting 100% in math tests, they are at the top of their class, etc. But what should the hotshot student think if after doing a math exercise we inform them that they have been slipped a reasoning-distorting drug, which wreaks havoc on people's mathematical abilities. Hearing this the student should probably lower their credence that they correctly understood the question and applied the right tools and consequently they should become far less confident that they got the right answer. But *before* they are informed about being drugged does their belief about their maths abilities look unreasonable? Of course not; it was the correct conclusion to draw based on their Higher Order Evidence up to the present moment. It is just that when they learnt that they had been drugged things changed.

The point is that the information about the reason-distorting drug is only the latest piece of evidence in a long series of evidence concerning their cognitive capacities. I hope it is becoming clear how similar this situation sounds to the cases of responding to first order evidence. In both situations you change your beliefs to reflect the information most recently gained, but in neither situation does it seem reasonable for you to have thought any differently before the most recent information came to light.

So it seems that responding to Higher Order Evidence doesn't commit the agent responding to any particular type of rational failure. But there are other candidate reasons for why one might consider Higher Order Evidence as specially interesting and different from first order evidence.

Candidate 2: Higher Order Evidence requires agent relativity

Christensen (2010) and Kelly (2010) propose another reason to think that Higher Order Evidence is importantly different from first order evidence. They propose that the evidential force of Higher Order Evidence is often relative to particular agents. Put simply, the impact a piece of Higher Order Evidence has will depend crucially on who the agent is.

To see this point consider the following example:

Don and Betty share some evidence E. Betty immediately comes to the conclusion that E is excellent evidence for P, and before Don has the opportunity to come to his own conclusion Betty shares her thoughts about E with Don. The question is; should Don consider Betty's opinion about E as *additional evidence* for P on top of the evidence E itself?

It seems that he should. If Betty is a reliable evaluator of evidence, her opinion is a good indication of how much support E confers on P. It's worth noting at this point that Betty's assessment of E doesn't confer support to P directly but indirectly. Betty's opinion confers support firstly on the proposition that 'E is excellent evidence for P' and if this is true, then P seems highly likely. I'll return to this detail later but for now I hope it's easy to see that the overall effect of Betty sharing her thoughts about E with Don is that Don will in the end become more confident of P. So Betty's assessment of E is, for Don, good additional evidence for P²¹.

The alternative is that Don fails to treat Betty's assessment of E as evidence for P in any way. This seems like a waste of useful information. As long as Betty is at least somewhat reliable why not utilise her views in order to come to the most accurate possible conclusion?

So the fact that 'Betty believes E is excellent evidence for P' is evidence for Don to become more confident of P. Things become more interesting when we consider whether or not Betty should also take the proposition that 'Betty is confident of P on the basis of E' as additional evidence for P on top of the initial evidence E.

It seems clearly not. Because if, like Don, Betty's assessment of E can be further evidence for P, then Betty could become increasingly more and more confident of P simply by learning that she has become more confident of P based on her evidence thus far. If Betty can become more confident of P based on her previous assessment of E then what is to stop her from becoming even still more confident in P based on her assessment of the Higher Order Evidence (i.e. the fact that 'Betty is confident of P on the basis of E'). This is the 'agent relativity' that Christensen and Kelly want to highlight; it seems that the Higher Order Evidence concerning Betty's assessment of E has evidential bearing for Don but has no such bearing for Betty.

Now in general, the idea that different people interpret the same evidence differently is not surprising. On Christmas morning when a family wakes to find that the milk and cookies they left for Santa have been consumed, the youngest children might

²¹ This point becomes especially clear when you consider that had Don never come to possess evidence E, then learning that 'Betty is confident of P on the basis of E' would likely be evidence for P insofar as Don believes Betty to be reliable and sincere.

understandably take it to be evidence that Santa visited them last night. Whereas their older and worldlier siblings would take the very same evidence to support the idea that their parents contrived to give the impression that Santa was here. In situations like this the different takes on the same evidence can be brought down to different background beliefs held by the agents involved.

What is striking about Higher Order Evidence is that agents can reasonably interpret the same evidence differently despite having the same set of background beliefs. In the case I've been discussing we can posit that Don and Betty are symmetric with respect to almost all their relevant background beliefs; they share both the first order evidence E and the Higher Order Evidence that 'Betty thinks E is excellent evidence for P'. And yet this second bit of evidence reasonably impacts Don very differently from how it should impact Betty. So Christensen argues: all that separates Don and Betty are indexical beliefs about who they are and therefore Higher Order Evidence impacts agents in an agent-relative manner.

Let's grant Christensen's point that the impact of Higher Order Evidence can vary between agents solely because of indexical beliefs. If this is true then an interesting question is: Why do indexical beliefs allow agent to respond differently to the same evidence? Christensen explains what is going on with Don and Betty by saying that the Higher Order Evidence offered in Betty's assessment of E can be used as a 'check' on Don's assessment (whether he has already made it or is about to make it). If Don is inclined to believe not-P on the basis of E then the fact that Betty has come to the opposite conclusion will cause Don to reconsider his assessment of E and raise his confidence in P. But Betty's assessment of E can't act as a check on itself; having judged that E is good evidence for P, Betty cannot, it seems, become even more confident that her assessment of E is correct simply by reiterating her previous assessment of E. As mentioned earlier, such a strategy would endorse Betty becoming maximally confident in her assessment of E (and maximally confident of P) based on conditioning on higher and higher orders of evidence.

This all seems fair enough, however Christensen's discussion doesn't explain further why Betty's assessment of E can serve as a check on Don's assessment of E and not fill this role for herself. Christensen is right that Higher Order Evidence can create the asymmetric 'checking' he describes, however what is more fundamental are the reasons why this asymmetry occurs. It is my contention that what we observe as the 'agent relativity' of Higher Order Evidence is in fact a far more general rule for responding to all evidence, whether it is First or Higher Order, and it's a rule that has little to do with who an agent is.

Candidate response 2: Agent relativity is no special part of Higher Order Evidence.

The general rule I am talking about I'll call the '*No Double Conditioning Rule*'. This simply states that you cannot, all other things being equal, condition on the same evidence twice.

The justification for this rule should be intuitive. Say you are planning an outdoor event on some weekend and you want to try and account for the chance of rain. You don't yet know where on the calendar the weekend will fall but you know that the frequency of rainy weekends is 30%. Using this information, you are 30% confident that it will rain on your chosen weekend. You then receive some evidence that this year there has been half as many rainy weekends thus far and there is strong indication that this trend will continue. Utilising this evidence you lower your confidence that it will rain on your event to around 15%. Now if you were to 'rediscover' the same evidence that indicated that there has been half as much rain as you would expect, it would be a big mistake to conditionalise and halve your confidence in rain yet again. At this point the evidence you rediscover is old evidence and as such it is already incorporated into your current confidence about your chosen weekend being a rainy one.

Apart from this intuitive justification for the 'no double counting' rule, a systematic case for this rule can also be made. One of the compelling results of Bayesian Epistemology is a clear explication as to how and why surprising evidence has such power to confirm a perhaps otherwise unlikely claim. In the famous Eddington Eclipse experiment, when light was observed to bend around the moon, this was powerful evidence for Einsteinian relativity, in part, because such a result was considered highly unlikely prior to the experiment. The flip side of this result is that the more expected some evidence is the less power it has to make significant revisions to our priors. The limiting case of this is certain evidence, which cannot shape credences in any way once that evidence comes in. According to orthodox Bayesianism any evidence that has already been conditionalised on, is certain. Indeed the act of conditionalising requires taking as certain whatever evidence is being conditioned on.

The proper conditioning procedure should go like this. Beginning with priors, we then encounter evidence, which we use to condition our existing prior. The result is a posterior belief that represents the sum of our information thus far and the newly generated posterior can become the prior for the next piece of new evidence we encounter. Central to this procedure is the proviso that *only new evidence is suitable to conditionalisation*.

Conditioning on exactly the same evidence more than once is the most drastic way to violate the '*No Double Conditioning Rule*'. More often the rule is violated in far more subtle ways by conditioning on some evidence which has already been partially

conditioned on when dealing with some other evidence. One very common cause of this is a failure to accurately assess independence / dependence relations between sets of evidence.

Say you have two diagnostic tests that test for the presence of a virus. Test1 is inexpensive and reasonably reliable, Test2 is more expensive but also more reliable. Your doctor is prudent and doesn't think it likely that you have the virus but in order to rule it out they order Test1 to be performed. When the results for Test1 come back positive this makes the prospect that you have the virus more likely, but since your doctor is so prudent they want to make sure and they order Test2. When Test2 comes back positive it seems the case is now very strong that you have the virus. Especially because when you come to conditioning on the result of Test2 your prior is already elevated by the positive result of Test1. But this would be a mistake. This is because Test2 is in fact a composite of Test1 together with other tweaks and alterations to make it overall more reliable than Test1, but not in a way that makes its results independent of the results from Test1. While its true that a positive result from Test2 sanctions a higher credence than a positive result from Test1, treating the two tests as independent would yield a more confidence than their make up warrants.

When judgements are independent they allow us to learn the most about whatever is being judged. When sources are not independent the amount we can learn from the various reports should be tempered by the extent that they are dependent.

With all this background covered we should now try and apply the notions of double counting and independence to a case of Higher Order Evidence and attempt to explain why Betty's assessment of evidence E has evidential influence over Don but not over Betty. When Don takes Betty's assessment of E to be additional evidence for P he can exploit Betty's opinion because up to that moment Betty's view is entirely new information about E's likely evidential bearing on P. Not only that but it is information which is independent from Don's judgement of E and so Don can legitimately utilise this information without the risk of double counting evidence.

Don can take his estimation of Betty's reliability and use it to calibrate his own assessment of E. Betty can't use her opinion to calibrate itself *in the same way* because it would likely involve drastic double counting. Betty's assessment of evidence E can be used by Don to calibrate his own assessment of E precisely because the two judgements are viewed as being independent. Whereas Betty cannot reasonably think that her assessment of E is independent from itself.

To summarise, Betty's judgement that 'E is good evidence for P' is clearly not independent from her previous conditioning on E. Whereas Betty's assessment of E *is* independent from Don conditioning on E. Acknowledging this it is clear why the Higher Order Evidence offered by Betty's assessment of E has evidential impact for Don and not for Betty. And notice particularly that explaining 'agent relativity' in this way doesn't involve any special role for Higher Order Evidence nor is the impact of evidence so crucially dependent on who the agents are.

While Christensen sees the apparent 'agent relativity' in response to Higher Order Evidence as proof that such is interestingly idiosyncratic, describing the differing response in terms of accurately gauging independence and avoiding double counting shows that there is nothing particularly special about Higher Order Evidence. First order evidence can't act as a check on itself. In exactly the same way, Higher Order Evidence cannot act as a check on itself.

The agent relativity seems like a special feature of Higher Order Evidence only because such evidence often has an indexical quality. But the indexical quality of a lot of Higher Order Evidence neither causes nor is uniquely correlated to the phenomenon described as 'agent relativity'. The true source of what Christensen is describing is the restriction on double counting evidence – a deep feature of responding to all types of evidence whether it is First Order, Higher Order, indexical or otherwise.

Are 'orders of evidence' intrinsic or relational?

I'll conclude this section focusing on Higher Order Evidence role in disagreement by emphasising something often overlooked by the commentators working in this area. Whether Higher Order Evidence is construed broadly as 'evidence about evidence' or more narrowly as 'evidence about someone's assessment of evidence' it is nearly always presented as an intrinsic property of the evidence under discussion.

When the medical intern is told they have been awake for 36 hours the higher order nature of that evidence is presented as an intrinsic aspect of that information because gaining this evidence occasions a change in their assessment of other evidence, namely their patient symptoms. This change in the assessment of the patient's symptoms results in the intern lowering their confidence in their diagnosis. This reasoning is presented as the reasonable, prudent and even rational thing for the intern to do; it is definitely the way I would want my doctor to respond if they were in the same situation. Ultimately we want our doctors to be mainly concerned about their patients and the accuracy of the diagnoses they make.

So seeing the evidence gained as higher order evidence, which ultimately bears on the patient's diagnosis, seems the correct way to take the evidence.

But I think the setup of this scenario, and similar ones discussed in the literature, bias us to view the Higher Order nature of some evidence as an intrinsic property of that evidence when in fact it isn't. Instead the first order or higher order nature of any evidence isn't an intrinsic feature of the evidence itself, but a label that describes a certain piece of evidence's relation to the particular propositions we are interested in. When we describe a given piece of evidence as Higher Order we are making a claim about what proposition we are primarily interested in.

To demonstrate this think about the example featuring Don and Betty. In that case Betty says that she thinks that 'E is good evidence for P' and this testimony is higher order evidence with respect to P. But what if we didn't care about the truth of P and were only interested in the truth of the claim that 'E is good evidence for P'. If Don wasn't concerned with the truth of P, but instead he was only concerned with whether 'E is good evidence for P' then Betty's assessment of E would be considered first order evidence. This is because Betty's assessment of E bears directly on the proposition that Don cares about. Generally, in cases of Higher Order Evidence the evidence received affects initially the probability of an intermediate proposition. It is the change in the probability of this intermediate proposition that then facilitates the change in the proposition we are most concerned with.

Now in more realistic examples we will care about the truth of many of the propositions that are affected directly and indirectly by some evidence. To capture this I suggest describing a given piece of evidence as having first order and higher order *aspects*. And depending on what proposition we care about most, either the evidence's First Order or Higher Order aspect will be most pertinent.

This point gives further support to my general argument that there is no interesting distinction between First Order and Higher Order evidence. If you accept the notion that the 'order' of some evidence is a relative property then the notion that Higher Order Evidence is particularly striking or special becomes even more difficult to motivate. At very least I think the burden of proof is squarely on anyone who champions to distinctiveness of Higher Order Evidence to make a better case.

Higher Order Evidence requires special inference rules

The final reason to consider why Disagreement might be special is that rationally responding to Disagreements might require special inference rules. This means that in order to take full advantage of the evidence presented when disagreements arise without stepping outside the bounds of rationality we need to restrict our existing methods of evidence evaluation with additional rules.

To see this, say I asked you to evaluate a simple math problem. You solve the problem easily, write down your answer and become highly confident that you have calculated it correctly. But I then inform you that just prior to being shown the math problem you were slipped a powerful drug, one that completely messes up the taker's mathematical capacities. Those who take this drug are completely unaware of its affects; they are blinded to the fact that they make simple and obvious errors in arithmetic. The nature of this drug is very specific, it doesn't cause visual hallucinations or widespread incompetence, it only hampers the takers ability to do the arithmetic. Someone under the influence of this drug can perfectly well spell out details of the problem,;what they can't do is reason from those details to reach the right answer. Put simply, the drug doesn't undermine evidence itself, it only undermines ones ability to correctly understand what that evidence supports²².

Almost everyone agrees that the intuitively reasonable result to draw from this example is that you should become far less confident in your initial answer. But having spelt out the details of this drug, both Christensen and Elga think that an additional restriction needs to be added in order to insure that the intuitively rational result is the one that our inference rules can deliver.

The apparent problem is that, as the setup makes clear, the evidence you currently possess directly entails one correct answer. And yet, given the evidence presented about the drug, you cannot draw the conclusion that is directly entailed by the evidence you possess. So, as Christensen says, responding to disagreements requires one to '*bracket off*' or '*set aside*' the conclusions that could ordinarily be drawn on the basis of that evidence. The alternative scenario would involve one's first order evidence remaining at full strength. This would mean that you would be reasonable in continuing to draw the conclusions entailed by your evidence about the arithmetic problem as you did prior to being told about the drug. If this were allowed, it's hard to see how learning about being drugged could have any effect on your confidence in your answer.

According to Christensen (2011) and Elga (2007), this problem is particularly striking when describing how to properly respond to cases of peer disagreement. Take the oft-discussed 'bill splitting case' where two people come to different conclusions on how much they owe

²² This example is adapted from a similar one found in Christensen (2011).

when splitting a dinner bill. Christensen is concerned that without the 'bracketing off restriction' there is nothing to prevent someone from reasoning like this:

My friend thinks our shares are \$25 each while I think that our shares are \$23 each. How should this information affect my confidence? Well, that will depend crucially on my friend's epistemic credentials: the more expert she is the stronger the indication that I have misinterpreted the relevant evidence about the bill. I normally think of her as being just as good a bill-splitter as me in situations like this, but it seems this time she got the wrong answer. My calculations tell me that the right answer is \$23, so she must have made a mistake and I don't need to lower significantly (if at all) my confidence in my initial result.

In this situation I have used the reasons for my own opinion to alter the epistemic credentials of my friend, and for Christensen doing this runs afoul of the 'Independence rule' which I raised on the introduction. This rule is a disagreement specific version of the bracketing off restriction; it states that if you disagree with someone about P you cannot use your reasons for P to come to a conclusion on the disputing parties epistemic credentials. Just like in the arithmetic example you need to 'bracket' or 'set aside' your reasons for your initial position and not let them influence your assessment of a disputant's epistemic credentials. In other words an epistemic evaluation of any disagreeing party should be 'independent' from one's reasoning about the issue under dispute.

Without these additional inference rules Christensen's and Elga's concern is that agents could reason in a way that appears, at least intuitively, circular. And indeed the idea of someone dismissing evidence on the basis of their reasoning when the reasoning itself is precisely what is likely to be undermined by the evidence they are dismissing, looks suspiciously circular.

However to make the case that the 'bracketing/independence' inference rules are not a special requirement of responding to disagreements, a few hurdles need to be cleared; Firstly, it needs to be shown that the circularity invoked in these cases is really a problem which needs to be solved. Secondly, if that hurdle is cleared it needs to be shown that the bracketing/independence rules can fix the threatened circularity problem. Thirdly, and most importantly, it should be shown that the bracketing/independence rules are particular to Higher Order Evidence and that they are not merely entailments of more basic inference rules already in place. I think that there is significant doubt over each of these claims.

The Bracketing/Independence inference rules may not be useable.

As already stated, Elga (2007) and Christensen introduce the 'bracketing/independence' inference restriction in order to prevent a kind of circular reasoning. But, as Christensen points out (2007, p.14) the kind of reason he is seeking to avoid doesn't beg the question in the normal way. When one disagrees with someone about P, it isn't straightforwardly circular to use one's reasons for P to downgrade the epistemic status of a disagreeing party. This procedure doesn't fit a classically circular argument structure or for that matter a rule-circular argument structure.

Given the reasoning procedure being discussed doesn't fit existing models of circular arguments we have two choices. One choice is to take the fact that the procedure doesn't fit into the existing framework as a sign that you've discovered a new and novel type of reasoning error; in this case it would be a special type of 'disagreement circularity'. The other choice is more eliminativist, and that is the option to reconsider whether or not the reasoning procedure is a type of circular reasoning error at all.

Christensen clearly opts for the former, while I'm far more attracted to the latter. For Christensen the fact that responding to Disagreement (and Higher Order Evidence more generally) threatens a unique type of reasoning error is yet more evidence of the specialness of Higher Order evidence. But in general if one suspects that some widget is an X and yet that widget doesn't conform to our current understanding of Xs, then I think there is a substantial burden of proof on those who claim they have discovered a new and novel type of X. A more parsimonious strategy would involve instead rejecting that the widget is an X and seeking to categorise the widget elsewhere. So while I certainly wouldn't rule out the possibility of 'disagreement circularity' being a genuine reasoning error, I think a stronger case needs to be made for it.

With respect to the second hurdle, I have doubts that the inference rules advocated by Christensen and Elga can adequately describe how to respond to Disagreements. This is because in order for the rules to be applicable to the range of relevant cases they need an unmotivated distinction.

In particular, both Christensen and Elga's restrictions need to maintain what appears to be a dubious distinction between one's *reason for a proposition* under dispute and *facts about one's reasons* for that same proposition. To bring out the distinction between reasons for the disputed proposition and facts about one's reasoning, Christensen describes a variation on his 'Bill-Splitting Case'. These examples involve disagreement by up-to-that-moment epistemic equals over the result of a basic arithmetic problem. In the original example,

these calculations are put down rather hastily and so the results of their mental math have a pretty large likelihood for error.

The variation I want to discuss Christensen calls 'Careful Checking'. In this scenario I am splitting a bill for lunch with a friend who I deem to be an epistemic peer. We are each spending a long time on the calculation, using pens and paper and even calculators. We then report our results and find a disagreement. He thinks we each owe \$25, and I think we each owe \$23.

What is importantly different in this variation is that the effect of my careful checking has the dual effect of raising my credence in my own answer whilst also lowering my credence in the answer of my friend once the disagreement is apparent. Like many commentators, Christensen thinks that it would be reasonable to explain this disagreement in terms of my friend's error as opposed to my own. But it is difficult to see what reasons independent of my reasons for a \$23-share could be used to favour an explanation in terms of my friend's error. But Christensen thinks I have such reasons, and in order to downgrade my friend's view I should reason thusly:

1. In arriving at my belief I use a highly reliable method.
2. The probability of two people using equally reliable methods of belief formation and disagreeing is very remote. The far more likely explanation is that one of us is less reliable than the other.
3. Which one of us is more likely to be unreliable? Well, I can more confidently rule out things that would make me less reliable (e.g. cognitive impairment or insincerity) for myself than for you. So I have reasons to think you are less reliable than myself.

Christensen thinks this reasoning is dispute-independent. He draws particular attention to the fact that it gives no specific support to the view that we each owe \$23; this reasoning could even have been followed before seeing the bill. (Christensen, 2011).

However both of these points do not yet establish that the reasoning just described is dispute-independent. Just because a reason doesn't expressly mention another proposition does not mean that the belief can't lend support to that proposition. The fact that I believe my vision is reliable is an important reason why, along with many other beliefs, I believe the traffic light is red or that thing on the roadside is a mailbox. My belief in the reliability of my vision says nothing directly about traffic lights or mailboxes, and yet if I didn't hold this belief I would be far less confident in both of these.

So to begin with I think that the distinction Christensen and Elga invoke is under motivated. But there are additional reasons for thinking that the distinction cannot do the work it is being called on to do. This is because, importantly, in order to accord with the Bracketing/Independence restriction, one's beliefs about their reasons for P clearly cannot also be reasons for P.

But are there beliefs about one's reasons for P that are not themselves reasons for P? Say I hear Scientists-X say that climate change is not caused by the actions of humans. In this example, my belief that 'Scientist-X told me so' is my reason for the belief that climate change isn't caused by human beings. There are certainly beliefs I might hold about this reason which are not themselves reasons for P. I could believe that when I heard Scientist-X speak she was wearing a green jacket. This may be, in some loose sense, a belief about my initial reason and this, by my own lights, has no bearing on whether or not P is true.

But it is clear that this kind of belief isn't going to do the work Christensen requires. What Christensen has in mind is something very specific. It is particularly my beliefs about my reliability that Christensen is interested in leveraging. So are there beliefs about my reliability that are not also beliefs in support of P.

What about the belief that 'Scientist-X is a respected scientist'? Is this a belief *about* the first reason for P or is it also a reason for P? I think it is a reason for P; if for no other reason than if I weren't to hold it I would be much less confident in P. Considering counterfactuals like this uncovers a vast set of beliefs we can think of as 'supporting reasons'. A supporting reason is one that gives detail to the force of the original reason it is connected to.

If I believe that it will rain today on the basis of dark clouds in the sky, it seems likely I will also have the supporting belief that dark clouds are highly indicative of rain. If I've found out the answer to a difficult sum by using a calculator, the strength of this belief will be the result of my credence that I entered the sum correctly and the calculator's reliability.

Returning to the climate change example, what about the belief, 'The reports of scientists-X are generally reliable'? This is also, I think, a supporting reason. What about the belief that 'the particular method of belief formation leading to this belief is highly reliable'? I can see no reason why this belief is different from the others. To my mind it is also a supporting belief. It is something that, although not immediately apparent from the initial reason, lies behind the initial belief as an *intrinsic part* of its epistemic force.

From Christensen's discussion it is clear he thinks one's confidence in his or her own method of belief formation is a belief *about* one's reasoning, not a direct part of one's

reasoning. For Christensen these two notions are distinct. But I don't think he has provided sufficient reasons to establish the distinctness of these notions. As mentioned before, even if a belief about one's belief formation process (e.g. I can reliably form beliefs on the basis of my ordinary vision) does not mention the content of some other belief, the Higher Order Belief might still be a reason for the first order belief. If the confidence you have in your belief formation process is a necessary part of the conclusion you come to (i.e. had you been more or less confident in your belief formation process you would have come to a different first order conclusion) then I think it deserves to be a genuine part of one's reasoning for the conclusion.

As well, Christensen's point that one could conduct reasoning about one's belief formation processes before actually putting them to use isn't conclusive either. I take it that Christensen thinks the following reasoning shows that one's beliefs about one's reliability are reasons for their reliability.

When the bill comes I will use a highly reliable method of belief formation. If my co-dinner disagrees with me, the chances of us both using equally reliable methods of belief formation are highly unlikely. Since I will have a greater knowledge of my own reliability, I can reasonably assume that it is me that is more reliable.

All that would occur in this instance would be a set of conditional statements which can be filled equally by one's reasons for the proposition despite not yet possessing them. As in:

When the bill comes I will have some set of reasons for my belief. If my co-diner disagrees with me, the chances that both of us have equally good reasons is highly unlikely. Since I have a greater understanding of my own reasons, I can justifiably assume that it is I that am most reliable.

So not only has Christensen not provided adequate reason to believe there is a distinction between one's reasons for P and one's considerations about one's belief formation method leading to belief in P, further investigation reveals there are strong reasons to doubt this distinction. In fact I think the very notion of a reason for P is inextricably tied to one's perceived reliability of the process leading to that belief.

Here is an illustration: if I buy a dozen eggs and use two to make an omelette, without counting I reason that there are ten eggs remaining in the carton. If someone were to ask me why I believe that there are ten eggs remaining I would explain my reasoning: that there were 12 and I used 2, $12 - 2 = 10$ and so 10 eggs remain.

The procedure just described would be, for Christensen, the reasons for my belief that there are 10 eggs remaining in the carton, in Elga's language it would be the detailed

‘chain of reasoning that led to (my) conclusion’ (2007, p.490) and this seems intuitively correct. It is these beliefs that, if someone were to disagree with me, I would need to set aside in evaluating their epistemic status. What I *could* use when judging someone’s epistemic status, according to Christensen, is my confidence in this method of belief formation. The fact that I consider this method of belief formation highly reliable allows me, according to Christensen, to reason thusly:

I followed what I believe is a highly reliable method of belief formation and have encountered a disagreement. It is extremely unlikely that two people employing equally reliable belief forming processes come to a disagreement; it is far more likely that one of us has a less reliable belief formation process. Which one of us is less reliable? Well I have more reasons to consider you less reliable than myself.

This is because I can more readily rule out many things which would decrease the reliability of my belief forming process. For example I can confidently rule out that I’m not drunk, tired or being insincere, but I cannot rule out these possibilities with as much confidence in your case. And so, because out of the two of us I think you are more likely to have followed a less reliable method, then I can justifiably dismiss your belief.

Christensen argues that this response to a disagreement concerning how many eggs are in the carton depends in no way on one’s reasons for believing that 10 eggs are in the carton. There is no mention of subtraction or the numbers 2, 10 and 12; there is not even a mention of eggs! And so it does seem that I haven’t relied on my reasons for the belief.

As mentioned earlier, I don’t think this is a compelling reason to think beliefs about the belief formation process leading to P aren’t themselves reasons for P. But if these notions are distinct, one should be able to conceive of them separately. That is, one should be able to conceive of my having the same reasons for 10 eggs being left in the carton, whilst I fail to have the beliefs about the reliability of the belief formation process leading to my belief that there are likely 10 eggs left.

What is crucial is that, when imagining this, it is important that the procedure identified earlier as my reasons for thinking 10 eggs are left continue to persist *as reasons for the end belief*. But I don’t think this is possible. If I were to enact the same mental procedure and yet I have no belief whatsoever as to the reliability of this procedure to yield correct beliefs, then the procedure would simply cease to be reason at all for the resultant belief.

But if these notions are distinct, one should be able to conceive of them separately. That is, one should be able to conceive of my having the same reasons for there being 10 eggs left in the carton whilst I fail to have the beliefs about the reliability of the belief formation

process leading to my belief that there are likely 10 eggs left. Again this seems impossible: if I have no view on the extent that these beliefs, when taken together, have a better chance of delivering a reliable belief, then why are they reasons? There is just as much reason to think that any other belief would be a reason for the proposition that there are 10 eggs left.

This discussion points to what I think is a simpler and more common-sense view of what beliefs constitute a reason for P:

To say that belief-S is a reason for P is just to say that belief-S is part of a set of beliefs that have the effect of raising one's credence in P and were they not to have belief-S then they would be less confident in P.

It seems to me that the reliability of one's method of belief formation is definitely a reason for one's belief. It is in fact the most important reason. What else is there to one's reasons for a particular belief other than the extent to which those reasons reliably support that belief?

These considerations, I think, show that Christensen's Independence restriction is at least under motivated or at worst deeply flawed. Not only is it not clear that the restriction does the work Christensen wants it to do, but I think there are also good reasons to think that the distinction it seeks to make use of is false.

Based on these observations I think the distinction between reasons for P and certain fact about those reasons, a distinction required by the Bracketing/independence Inference rules, is untenable without further argument.

The third problem is that Bracketing/independence rules apply uniquely to the treatment of Higher Order Evidence and they are genuine additions to the familiar set of inference rules. In this instance I am thinking specifically about the set of inference rules which come directly from the Bayesian framework. These are conditionalisation and all its consequences, some of which I have already discussed like the 'no double counting rule' and basic probabilistic independence.

Conclusion:

Returning to the purpose of this paper, I have attempted to show that based on the literature so far, disagreement is not special. By casting disagreement in a Bayesian context I have attempted to show that cases of disagreement involve neither responding to special evidence nor do they require the use of special inference rules.

However, this is not to say that disagreement, and in particular a Bayesian understanding of disagreement, isn't philosophically interesting and worthy of investigation.

Bayesianism is a controversial and innovative epistemological framework, with many ideas and concepts that are yet to be applied to cases of disagreement. In the next chapter I will attempt to bring some other tools from Bayesian epistemology and put these to work on cases of disagreement in a bid uncover unattended to aspects of disagreement.

4. Disagreement and 'Belief Independence'.

When should numbers matter?

Walking out of a movie one night I turn to my friend and say; 'Wow, that Kevin Bacon sure can act. I thought he was very convincing in this movie'. Sounding a little surprised, my friend responds, 'I think you're confused - Kevin Bacon wasn't in the movie we just saw'. 'That's strange', I think to myself. I really thought that was Kevin Bacon in the movie, but upon hearing that my friend disagrees I become a bit less sure.

What if I'd seen the film with a large group of friends, all of whom disagreed with me about Kevin Bacon being in the movie? If that had occurred then I would be *a lot* less sure; I would lower my confidence substantially more than I did in the single disagreeing party case above.

This story is an illustration of a strong intuition many have about disagreements: A disagreement with a group carries more epistemic force than a disagreement with just one lone disputant. In short: when responding to disagreements, numbers matter.

In general, discussants within the epistemology of disagreement agree this intuition is correct. But many (Elga (2010), Kelly (2005), Goldman (2001), McGrath (2008)) think that it is limited in an important way. In particular, they hold that there are some cases where numbers decidedly do not matter. These are cases where the individuals within a disagreeing group are not formed independently. Elga sums up the idea nicely here:

'An additional outside opinion should move one only to the extent that one counts it as independent from opinions one has already taken into account.' (Elga, 2010, p.3)

To pick up the cinema example, consider what difference it would make if there was one person in the group who was dominant and whose opinion instantly swayed everyone else's once they made their thoughts known. If I knew such a person was in the group and that they disagreed with me, it would come as no surprise that the entire group disagreed with me. In light of this information, once I had taken into account the opinion of my dominant friend it would be wrong for me to go on and further alter my confidence that Kevin Bacon was in the film due to the group disagreeing with me en masse.

Introducing Lackey and Belief Independence

In her paper 'Disagreement and Belief Dependence', Jennifer Lackey identifies the thesis behind this common intuition and calls it 'Belief Independence'.

Belief Independence: That for someone to be influenced by the beliefs of an epistemic peer, the epistemic peer's beliefs must be independent from views already considered.

Belief Independence states that for a disagreeing peer's belief to rationally sway another, that belief must be independent from views already considered by the agent being swayed. So in a scenario involving three peers - A, B and C – where agent A has already encountered and responded to a disagreement with agent B and then agent A discovers a disagreement with agent C, the opinion of agent C can only have epistemic force if agent C's belief is independent from the belief of agent B which agent A has already taken into account.

A typical justifying example might be responding to a rumour you heard from two people, call them Simon and Chris. Hearing the rumour initially from Simon gave it relatively little plausibility, but hearing the same rumour from Chris a few days later seems to give the rumour some additional support. However if you were to discover that Chris only believed in the rumour because he heard it from Simon, you would probably be less confident that the rumour is true than you had been when you thought that Simon and Chris had formed their views independently.

It is worth noting that in actual cases of rumour spreading as well as many other real-world disagreements, the involved parties would likely not be *completely* dependent or independent. There is a lot of room between these extremes for what I'll call 'partial dependence' in the rest of this chapter.

The limiting case would be one where the second party's views were completely dependent on the first's – where the second party exercises no epistemic agency and acts simply as a mouthpiece for the first party's views. In such a scenario of total dependence it seems clear to everyone that the view of the second party offers no additional epistemic force beyond those present in the view of the first party.

But despite its popularity and intuitive appeal, Belief Independence has received very little direct argument in favour of it. It appears to be an idea which is so accepted that very few scholars have bothered to explicitly argue for its truth. In this context, it is understandable and even laudable that Lackey sets out to investigate and eventually problematise it. However, a close analysis reveals that Lackey's argument against Belief Independence is deeply flawed.

In showing how her argument fails, we can better understand how we should respond to disagreement. In what follows I will mount a defence for Belief Independence. I will start

by closely examining Lackey's criticisms of Belief Independence. I will then go on to explain how Belief Independence fits within a larger epistemological framework, and in so doing shed light on why other commentators haven't offered explicit arguments in favour of it in the past.

Lackey's argument against Belief Independence.

In a nutshell, Lackey's argument against the Belief Independence thesis is that there is no interpretation of the thesis that makes it true in all cases.

In her paper, Lackey discusses a number of different forms of the Belief Independence thesis, each stemming from a slightly different conception of what it means for agents to be dependent/independent. The interpretations of Belief Independence which Lackey discusses are 'likemindedness', 'source dependence' and 'conditional dependence'. For each of these she then proposes arguments and counterexamples which she says show that Belief Independence is false.

To examine her argument I'll go through each of these interpretations of dependence/independence, starting briefly with 'likemindedness'.

Likemindedness

This is simply the idea of using likemindedness of individuals as a proxy for dependence. Understood this way, individuals who are significantly likeminded are dependent. On the opposite end of the spectrum, individuals who aren't at all likeminded are independent.

It seems pretty clear that this isn't a good way to think about belief independence for two key reasons. Firstly, it doesn't acknowledge that agents can arrive at similar beliefs in very different ways. These types of views are precisely the ones that can be epistemically valuable – think of two scientific studies coming to the same conclusion using completely different methodologies. Secondly, it makes belief independence far too strong. It could render rational agents unresponsive to the views of anyone whose opinions are similar to those they have already taken into account, creating a 'heard one, heard 'em all' type of epistemology.

Source Dependence

Source Dependence has more going for it as an interpretation of Belief Independence. The idea here is that individuals' beliefs are dependent on each other when they are dependent

on the same source. Of course, in many cases two or more people will share some evidence but they will also have some evidence unique to them; they can be said to be *partially source dependent*. But in this instance we are interested in the more extreme cases, where agents are *wholly source dependent*. Using this interpretation, Lackey formulates a new version of Belief Independence:

Belief Independence 2: When A disagrees with peers B, C and so on with respect to a given question and A has already rationally taken into account the disagreement with B, A's disagreement with C and so on requires doxastic revision for A only if the beliefs of C and so on, are not wholly dependent on B's belief. (Lackey, 2013)

With this interpretation Lackey begins her main strategy for arguing against Belief Independence using a series of what she believes are counterexamples, arguing that agents can be wholly source dependent and yet still have beliefs that carry their own distinct epistemic weight.

Here is one of Lackey's proposed counterexamples: Lackey asks us to imagine a scenario where a blood spatter expert called Cathy, is asked to give her opinion on whether a particular defendant is guilty based on the blood spatter evidence at a crime scene. After coming to a conclusion, she shares her opinion with two other people, Abby and Betsy. They are equally knowledgeable about the science of blood spatter analysis, but crucially they do not have access to the original evidence Cathy used to form her opinion. They only have access to Cathy's testimony that the defendant is guilty.

But despite being wholly source dependent on Cathy's testimony, Abby and Betsy are not epistemically symmetrical. It turns out that Abby has the capacity to 'autonomously filter' Cathy's testimony by 'critically assessing its reliability, monitoring for defeaters and comparing the content of belief that she forms with her background beliefs' (Lackey, 2013, p.). Betsy, on the other hand, has no such filter. And so, claims Lackey, a peer that disagreed with Abby and Betsy would be rationally required to change their mind more than they would were Betsy the lone disputant.

The other examples Lackey describes have the same basic structure. In each case, Lackey's general strategy is to make the agents in her examples dependent but still allow them to exhibit some agency in forming their views. This allows the views of these agents to have some commensurate epistemic force beyond that of the views of those they are largely dependent on.

Thus Lackey claims that she has described a scenario where agents are wholly dependent (in this case source dependent) and yet the views of the many still carry more epistemic weight than the views of a single dissenting party.

However, there are two key reasons why Lackey's example does not truly meet this description, as I will explain below. Firstly, Abby and Betsy shouldn't be considered peers. Secondly, Abby and Betsy are not wholly source dependent.

Abby has evidence Betsy lacks. Lackey says that even in scenarios where Abby's belief is based on Cathy's opinion, Abby may still exercise some epistemic autonomy. She can do this by applying beliefs that she has about Cathy's reliability, in general or in this specific instance. This is interesting, because to my mind this factor means that Abby and Betsy are not epistemic peers. Whatever information Abby is using to appraise Cathy's opinion would also make Betsy more confident of her own opinion if she possessed it.

In Lackey's example Abby is a better judge of whether Cathy's testimony is in fact correct. The straightforward reason for this is that Abby has some evidence Betsy lacks, namely that testimonies from sources like Cathy are often very reliable. This would mean that Abby and Betsy *aren't evidential equals and nor they are epistemic peers*. So it isn't surprising that each has some unique epistemic impact.

To see this point another way, consider what might happen if Abby and Betsy are repeatedly called upon by the courts to act as corroborating expert witnesses. Just as in the original example, in each of these cases they receive the identical testimony from an expert witness (Cathy₁, Cathy₂, Cathy₃ ... Cathy_n), each of which has the same high epistemic status as the original Cathy. As before, Abby and Betsy need to then decide if they agree (or to what extent they agree) with each Cathy_n.

Say they do this 10,000 times. While we wouldn't expect Abby and Betsy to agree on all cases, we would expect that as they do more and more cases the number of cases where they adopt the opinion of the expert witness would converge. But from the information Lackey has given us, this wouldn't happen because Abby is just more sensitive to errors made by the expert witnesses than Betsy is. So over time Betsy is going to overlook mistakes the Cathies make that Abby might pick up on.

If we imagine that at some point the true accuracy of each expert witness was revealed and we could assess the track records of Abby and Betsy, they would reveal that Abby is simply a more reliable judge of blood spatter experts than Betsy. This quality disqualifies Abby and Betsy from being peers in the domain of blood spatter experts.

So what Lackey has described is just a case where two people receive the same testimony but one person is better equipped to make good use of that testimony. Hence the opinions that person forms on the testimony's basis are more informative.

Lackey anticipates this objection and responds by saying that in her example Abby's background beliefs do not *ground* the beliefs arrived at, rather they *act as a filtering device* and so whatever evidence Abby utilises in this filtering doesn't affect her status as being wholly source dependent in the same manner as Betsy is.

The obvious question is: what is the difference between something 'grounding' a belief versus it 'filtering' a belief? This is a question I will come back to in the section labelled 'what is evidential support?', but first I would like to introduce the last interpretation of dependence that Lackey considers.

Conditional Dependence

Here Lackey follows work by Goldman (2001) to use conditioning as a way to provide an account for belief dependence. According to Goldman, Y's belief in H, notated as Y(H), is 'non-independent' if their belief fits the following:

Non-independence: The chances of Y believing H are the same, conditional on X believing H and H being true as they are when X believes H and H is false.

$$P(Y(H) \mid X(H) \ \& \ H) = P(Y(H) \mid X(H) \ \& \ -H)$$

Visualise this as Y forming their belief on that basis of X's belief and nothing else, not taking any time to examine the truth of H for themselves. Following Goldman, Lackey says that for individuals Y and X, Y is a 'non-discriminating reflector' of X's belief in H when Y satisfies the non-independence relation.

Using this idea Lackey constructs another interpretation of Belief Independence, called Belief Independence³. This states that for additional disagreeing peers to have epistemic force those peers cannot be non-independent from disagreements already taken into account.

Lackey attacks Belief Independence³ using a similar technique to one we have seen before. Returning to her former example, this time Abby and another blood spatter expert called Annie are both non-discriminating reflectors of Cathy's opinion. As such they are just as likely to believe Cathy's opinion, regardless of its truth value.

However, while Abby and Annie are both non-discriminating reflectors of Cathy's opinion of H, they are still epistemically distinct. It turns out that Abby (but not Annie) is highly discriminating when it comes to evaluating the sources on which she basis her opinion of H. Annie, on the other hand, is completely indiscriminating and she adopts Cathy's opinion of H without exercising any particular judgement about Cathy's ability to correctly evaluate H.

This, according to Lackey, means that Abby can do a similar trick to what we saw earlier and filter Cathy's testimony. This filtered testimony means that Abby's opinion has additional worth beyond Annie's. This would necessitate more revision from a peer who disagreed with both of them, were one to arise.

This argument is wrong for two reasons. Firstly, it doesn't fit Goldman's account of non-independence, so Lackey hasn't described a counter example to Belief Independence³. And secondly, it doesn't have a coherent account of how evidential support operates. Let's examine both of these.

Problems with epistemic 'filters' in non-independence

In Lackey's previous two examples, Abby has certain background beliefs about the reliability of Cathy with respect to the proposition in question (the defendant's guilt). Abby is also able to screen for defeaters when she learns about Cathy's opinion and so, Lackey claims, she is able to act as an 'epistemic filter' for Cathy's testimony.

Generally speaking, filters work by letting through what we want to let through and catching the stuff we don't, leaving us free to dispose of it. It's an appealing analogy for the selection process involved in responding to evidence. Of course, filters don't work perfectly, they don't take out 100% of the bad stuff and inevitably they will also select out some of the stuff we'd like to hang on to. And as it is with real life filters, so too with epistemic ones, for if we possessed perfect filters we would never be fooled by evidence. We would only hang on to the true propositions and we would only let go of the false ones.

If Abby is in any sense filtering Cathy's testimony about H, then there must be some operation on that testimony. Lackey offers up examples of the types of operation she performs when she filters; screening for defeaters, checking her existing views on Cathy's trustworthiness, etc (Lackey, 2013, p.259).

When Abby screens for defeaters, in order for that filter to have any effect there must be a non-zero chance that she can find and then act on the defeaters she discovers. And wouldn't the likelihood of finding such a defeater be affected by the actual truth or falsity of H? If not, that is if Abby's ability to screen for defeaters had *no relationship to the truth of H* then she would be an entirely useless epistemic filter.

On the other hand if Abby is more likely to believe H on the basis of Cathy's opinion when H is true than when H is false, then Abby does not meet the criteria for Goldman's definition of non-independence. Abby is at least partially conditionally dependent.

So Lackey has not shown by this example that it is possible for someone's view to be completely conditionally dependent and yet still carry their own epistemic force in disagreements.

What is evidential support?

The other problem with Lackey's argument against Belief Independence³ harks back to the problem with the argument against the source dependence account. In that case Abby utilised evidence (about potential defeaters or reliability) to filter Cathy's testimony, but why isn't this evidence considered part of the source evidence in favour of Cathy testimony? Likewise in the case of conditional non-independence, why isn't Abby's evidence about Cathy's reliability itself evidence for Cathy's testimony?

As mentioned before, for the source dependence interpretation, Lackey says evidence used as an epistemic filter for some other proposition is not evidence for the filtered proposition (that the defendant is guilty) because it is not *grounded* in the evidence used in filtering. But how can we be sure that what Abby comes to believe is not grounded in the filtering evidence? What is the test for deciding if one particular belief is grounded in any particular piece of evidence?

Unfortunately Lackey gives us very few clues on this. She does make the point that the grounded relation is not straightforwardly causal (Lackey 2013, p.247). So we can rule out a causal account which just says that 'belief B is as grounded evidence E just when E causes B'. But eliminating this option doesn't bring us much closer to knowing what being evidentially grounded actually means.

So let's look elsewhere, to the Bayesian account of evidence. One of the very appealing features of Bayesian epistemology is its intuitive and powerful account of evidential support.

For a Bayesian, for any hypothesis H and any evidence E, E is evidence for H if, and only if, the following is true.

$$\Pr(H \mid E) > \Pr(H)$$

This means that E is evidence for H if the probability (or credence for Subjective Bayesians) of H given E is greater than the probability of H on its own. There is much to say (and much to debate) about how precisely to understand this concept of evidence, but for my purposes here I'd like to draw attention to a couple of its features that many find striking.

Firstly, this account applies equally well to situations where we have evidence E and those where we are yet to discover E. All that is required is for us to come to a conclusion, based on our beliefs so far, as to how the attainment of E would change the chance of H occurring.

Secondly, this is a process which is often internal to the agent examining the evidence, meaning that the same evidence can legitimately and rationally confer different amounts of support on H. This means the Bayesian account is potentially highly subjective and this can worry some. However, in this case I think it is a feature and not a flaw because it seems to gel well with Lackey's description of how two people can diverge even when responding to the same evidence.

Lastly, the Bayesian account of evidence is a helpful framework because its formal notation can make it clearer exactly how evidence confers support on particular propositions.

Let's apply this framework to the type of cases Lackey describes. According to the Bayesian account of evidential support, does evidence for someone's reliability count as evidence in favour of the content of their testimony?

One way of show that it is would be demonstrating that such evidence is transitive. If evidential support is transitive then if A is evidence for B and B is evidence for C then, by transitivity, A is evidence for C. But unfortunately evidential support is not generally transitive; consider the following example. You have just done a very survey of technology professionals and in the results indicate two major findings: that people who work for large technology companies are more likely than average to know how to code and the majority of people who know how to code uses Mac computers at their workplaces. In

response to this Google decided to ban Mac computers from their workplaces. With this information as background, consider the following propositions:

Q: X works for Google.

R: X knows how to code.

S: X uses a Mac at work.

From the information above we can say that Q is good evidence for R and that R is good evidence for S. But we also know that Q is not good evidence for S (in fact Q disqualifies S from being true) because of Google's decision to ban Mac computers at work. So evidential support fails to be transitive in this case.

Transitivity and probabilistic support

However, there are some cases where probabilistic support is transitive, and cases of dependent disagreement and testimony, the very scenarios Lackey describes, might be just such cases. If this turns out to be the case, it would show that Lackey is incorrect in claiming that evidence for a testifier's reliability isn't itself evidence for the content of their testimony.

Tomoji Shogenji (2003) shows that probabilistic support is transitive between three propositions when the intermediate proposition 'screens off' support from the initial proposition to the final proposition. Or, put another way, once the truth or falsity of the intermediate proposition is established, the truth of the initial proposition no longer holds any sway over the probability that the final proposition is true.

The 'screen off' condition, where B screens off A with respect to C is as follows:

$$[\Pr(C \mid B \ \& \ A) = \Pr(C \mid B)] \text{ and } [\Pr(C \mid \neg B \ \& \ A) = \Pr(C \mid \neg B)]$$

Notice that this screen off condition fails in the Google example above because the probability that someone uses a Mac at work is different in the case that they work for Google and know how to code than in the case where they just know how to code. Put differently, it fails the identities in both conjuncts of the 'screen off' condition.

Now let's see this in the context of the example we have been discussing. Remember in the example Cathy is an expert who believes and has publicly testified that P and Abby has access to some evidence about Cathy's reliability (though screening for defeaters or other means). Now consider the following three propositions:

A: Abby remembers evidence that Cathy is reliable.

B: Cathy is reliable

C: P is true.

Plugging in our three example propositions, A, B and C, we can see that the screen off condition does hold, and the second proposition screens off support from A to C. The first statement says that once we have already established that Cathy is reliable (proposition B) then the fact that Abby remembers evidence of Cathy's reliability has no bearing on the truth of C. The second condition makes a similar claim, that if we establish that Cathy is not reliable (-B) then the fact that Abby remembers evidence that she is reliable should have no impact on the truth of P. In short, if you have established someone's reliability, then you should also probably believe that Abby is mistaken in thinking they are guilty, regardless of how many people testify otherwise.

The result is in keeping with Shogenji's general finding that 'probabilistic support is transitive provided the original evidence is testimonial, memorial or perceptual, and the intermediary proposition is its representational content' (Shogenji, 2003, p.615). And it appears that many of the cases which involve judgements about someone else reliability will involve starting from testimonial, memorial or perceptual evidence. Therefore, I think, we have additional grounds for saying that if an Abby comes to believe P on the basis of Cathy's testimony, then Abby's belief in P is 'grounded' (Lackey's terminology) in Abby's evidence about Cathy's testimony. This will be the case for whenever Shogenji's 'screen off' condition holds and while we can't say that it will hold on all cases of disagreement, it will hold in many as they often involve evidence which is testimonial, memorial or perceptual.

Lackey's response

When responding to each of Lackey's arguments, I have argued that she hasn't succeeded in providing genuine counterexamples to Belief Independence because in each case the actors involved are not truly dependent.

Belief Independence states that only disagreements with two (or more) non-dependent agents count as more than one disagreement. However, I have shown that Lackey has not succeeded so far in describing such a counterexample.

So there we are, case closed, Belief Independence is saved. Well, not quite.

Lackey anticipates this charge and responds by pointing out that in cases of complete dependence the agents involved will not be epistemic peers and seeing as Belief Independence is a thesis concerned *only* with disagreements among epistemic peers, such cases, argues Lackey, are not relevant to a discussion about Belief Independence.

According to Lackey, completely dependent agents cannot be peers; if Colin is truly a completely blind follower of Julie's opinion on a certain topic, no matter how Julie arrived at her view, then it is difficult to see how Colin's epistemic status could equal Julie's. If Colin is simply parroting Julie's views, it doesn't seem likely that he is employing equal evidence or cognitive ability with respect to the issue at hand.

The place of epistemic peerhood in Belief Independence

Here is where, I believe, we can find a fundamental oversight in Lackey's approach; like many in the epistemology of disagreement, Lackey seems to be very caught up in the notion of peer disagreement. But in fact there is no good reason to think that the Belief Independence thesis applies only to cases involving disagreeing peers. On the contrary, it should be a general principle which applies to peers and non-peers alike; it can be seen as a general feature of all disagreements, or even more broadly, as a feature of all evidence.

To motivate this idea a little, here is a simple example of disagreement among non-peers where Belief Independence seems useful. Take two people who I think are epistemic peers with respect to each other and epistemic superiors with respect to myself. If these epistemic superiors were completely dependent on each other then discovering that they both disagree with me offers me no more evidence than if I discovered that one of them disagreed with me. In this case, the fact that the two people who disagree with me are my epistemic superiors has no bearing on the applicability of Belief independence.

From examples like these it seems at least initially plausible that something like Belief Independence might apply to cases of non-peer disagreement. So why does Lackey insist that Belief Independence only applies to peer disagreements?

This question is hard to answer, in part because Lackey doesn't give a detailed account of her standards for epistemic peerhood. She defines epistemic peerhood in terms of evidential and cognitive equality, specifying that two agents are epistemic peers if they are broadly equal in their acquaintance with the relevant evidence for a proposition and equal in the cognitive abilities brought to bear on that same proposition²³. But this doesn't really give us any more clues as to why Lackey thinks Belief Independence only applies to peer disagreements. To start to get an idea, I think we have to look a bit more broadly at how the discussion has evolved within the epistemology of disagreement.

It is certainly true to say that up to this point, the epistemology of disagreement has been primarily concerned with how to rationally respond to peer disagreement. As I have argued in a previous chapter, the reason for this focus stems from peer disagreement being an 'edge case' where different people's intuitions diverge as to how a rational agent would respond.

The tendency for a general philosophical topic to become focused on an edge case (or set of edge cases) is far from novel. In modern philosophy it has become common. Think, for example, of the discussion of Gettier cases in the epistemology of knowledge; fission cases in discussions of personal identity; or philosophical zombies in the philosophy of mind. Each of these are extreme examples around which debate centres not because they are typical of our actual experience of knowledge, personal identity or minds, but because they defy some of the most influential accounts of those concepts.

The reason for discussing Gettier cases isn't because we are interested in the intrinsic features of Gettier cases on their own. We are interested in all of these edge cases because they reflect on our understanding of some larger concept. Gettier cases are interesting because they show us something about knowledge, and when I say 'knowledge' I mean it in a broad sense. It is the very thing that we're speaking about when discussing that large range of things we claim to know.

It seems that some (and specifically Lackey in this instance) have forgotten this point. To put it plainly; we should be interested in peer disagreement not only to work out the specific features of these cases or because different philosophers have different intuitions about how one should respond to a peer disagreement. We are interested in cases of peer

²³ Despite the problems with this definition, this is the dominant way of understanding epistemic peerhood. I discuss it in detail earlier on as the 'majority view'.

disagreement because the same concepts are at play in all disagreements, peer and non-peer alike. As I have said previously, issues raised in peer disagreements, like evidence, reliability and reasons are some of the most important notions in the whole of epistemology.

I would therefore make the claim that just as an account of peer disagreement reflects on all disagreements, so too accounts of disagreement impact directly on our general understanding of how evidence works and how a rational agent should respond to it.

Given these observations, I think we should not assume that Belief Independence only applies to cases of peer disagreement; to do so would just be a quirk of philosophical methodology. There are grounds to think that Belief Independence should have a much wider application.

The ‘Information Interpretation’ of Belief Independence

Now that I have addressed Lackey’s explicit arguments against Belief Independence I would like to end by proposing an account of it which, I think, puts the principle in a more appropriate context. I call it the ‘Information Interpretation’ of Belief Independence. This is as follows:

When A disagrees with peers B, C and so on with respect to a given question, and A has already rationally taken into account the disagreement with B, A’s disagreement with C and so on requires doxastic revision for A insofar as the beliefs of C and so on represent new information for A.

This interpretation of Belief Independence in the context of peer disagreement is an instantiation of a far more general principle concerning all evidence.

Evidence Independence: All evidence holds epistemic sway over agent A insofar as that evidence represents new information for A.

According to the Information Interpretation, evidential (or epistemic) independence is very similar to probabilistic independence. Probabilistically speaking, two events are independent if the outcome of one has no impact on the chances of the other. Similarly, two pieces of evidence are epistemically independent if learning about one provides us with no additional information about the other.

So why do we need this new interpretation, particularly as I have been arguing that Lackey’s arguments against the interpretations she discusses do not succeed? The answer is that while the interpretations Lackey raises aren’t flawed in the way that she thinks they

are they, they are less than complete. For example, interpreting Belief Independence as source dependence required Lackey to employ a notion of ‘evidential grounding’, where two (or more) agents’ beliefs were dependant when they were grounded in the same evidence. But this interpretation is still missing an account of what it means for a belief to being ‘grounded’ in some evidence²⁴. What I am pointing to here is not necessarily a fatal flaw, but a gap which could be filled in. But I think we have one additional reason to not bother filling this gap and instead embrace the Information Interpretation of belief Independence.

The Information Interpretation clarifies the underlying motivation behind Belief Independence and that is to prevent us from double counting evidence. This is the fundamental reason why ideas very similar to Belief Independence keep popping up in the epistemology of disagreement²⁵; counting evidence you already have as novel is inductively irrational. This notion also has strong intuitive appeal, if it were permissible for exactly the same evidence one already had to have additional epistemic impact then agents could perhaps become arbitrarily confident of any proposition that one had at least some positive evidence in favour of it.

But we need not be satisfied with this intuitive argument for the irrationality of double counting evidence because it can be further demonstrated by a diachronic version of the well known Dutch Book Argument (Teller 1973 and Lewis 1999). Dutch Book arguments seek to demonstrate the features of rational epistemic behaviour via a pragmatic self defeat test. They do this by showing what conditions an epistemic agent would have to satisfy to be vulnerable to a particularly devious bookie – someone who offers you a series of bets that you accept as fair but would also guarantee you a sure loss. In short they show that if epistemic agents’ beliefs have some feature then acting on the basis of that belief would mean that they lose no matter what the outcome. It is then proposed that any feature which allows this type of ‘pragmatic self defeat’ is irrational.

The conclusion of this type of argument is that if an agent has beliefs which violate any of the axioms of probability then they are vulnerable to accepting a series of bets that guarantee them a net loss. If utilising your beliefs in order to pursue your goals means that you are guaranteed to end up frustrating those goals, then it seems like something has clearly gone very wrong with your beliefs.

²⁴ The Information Interpretation is consistent with the Bayesian specification of ‘evidential grounding’ is propose in the section ‘what is evidential support?’.

²⁵ I have already discussed a number of similar cases in this thesis in the context of Higher Order Evidence and the special inference rules of disagreement.

Dutch book arguments attempt to describe the limits of validity for inductive inference in a similar way to how validity can be argued for in the context of deductive inference. In deductive rationality having true premises and reasoning about them with valid inferences will always yield true conclusions. If we are certain that we have true premises and reasoned about them using what we think are valid inferences, then we expect the conclusion we draw to be true. But if the result of that reasoning yields a false conclusion, then we have to admit that something has gone wrong in our reasoning. We have discovered that the inferences we were using are not valid.

Something similar is happening when dutch book arguments are used to describe inductive rationality. But the aims of deductive and inductive validity are slightly different. In a deductive context, validity preserves truth: if an argument is valid it is certain that true premises will yield a true conclusion. But this is not the case in inductive validity, where one can take a corpus of evidence, understand it correctly, make a prediction based on it and still get the prediction wrong. Sometimes evidence misleads us or we are just unlucky.

So instead of being truth preserving, what is the distinguishing feature of valid inductive reasoning? The answer is to *guide action in an uncertain world*. And, as the dutch book arguments demonstrate, credences which are susceptible to a Dutch Book will be a terrible guide because such beliefs can fail in a completely new way by actively undermining each other.

Returning to the case of double counting evidence, we can surmise this behaviour with the following, where Cr_1 represents an agent's credences at one time and Cr_2 represents an agent's credences at a later time.

$$Cr_1(H \mid E) \neq Cr_2(H \mid E)$$

This simply states that at two different times, but conditioning on the same evidence, the agent has a different credence about the truth of H . This is a violation of conditioning and according to the diachronic Dutch Book argument (Teller, 1973 and Lewis 1999), any agent whose beliefs match the statement above would be susceptible to a Dutch Book. So double counting evidence is inductively irrational and treating old evidence like new evidence is an instance of double counting. For these reasons I think the Information Interpretation is a superior treatment of the Belief Independence principle.

But the advantages for the Information Interpretation go further: it allows us to quantify exactly how much information a particular piece of evidence represents. This is one of the great features of a Bayesian approach to epistemology versus a traditional all or nothing

model of belief. Speaking broadly we can think of evidential relationships as falling into three categories; independent, dependent and completely dependent. The issue is that an all or nothing view of belief can make these categories seem somewhat equal in size, but this is not the case. The category of 'dependent' evidence is infinitely varied - from very weakly dependent to very strongly dependent - whereas the categories of independence and complete dependence are singular and extreme.

An Information Interpretation allows the level of dependence among information sources to influence the amount of information coming from those sources. It does this by characterising the amount of information a belief represents as a measure of how surprising that opinion is given our prior beliefs up to that moment. This understanding of the quantity of information is the famous definition of information in Claude Shannon's 1948 paper 'A Mathematical Theory of Communication'. One of the features of Shannon's definition for us is that it can be used to describe the amount of information of individual messages and the expected information output of a system, also called the system's 'entropy'. We can then think of groups of people and their beliefs – which could be independent or dependant – as information systems.

If a group of people's views are completely dependent, discovering what each member of the group thinks one by one isn't giving me any additional information, because knowing the opinion of any one member allows me to perfectly predict the opinion of all members. This scenario contrasts with a group of people whose opinions are highly dependent (but not completely dependent) on each other. When I learn that one member of the group disagrees with me this will change my belief about the relevant proposition, reflecting my view of their reliability on that proposition. But if I then learn the opinion of another member of that same highly dependent group, this still has the capacity to change my credences because the second person's opinion is still uncertain to me. But interestingly, provided the opinion of this second partially dependent person agrees with the first, the second person's opinion will likely have a smaller effect on me because their belief was unsurprising given the opinion of the first group member.

Conclusion

To conclude, the intuition that many in the epistemology of disagreement had seems to be validated: that when responding to multiple cases of disagreement on the same subject, each disagreeing individual will carry epistemic weight insofar as those agents are

independent. But although the principle of Belief Independence is vindicated, it is so not as a special feature of disagreement but as yet another feature of all evidence. Belief Independence – in all contexts – rephrases the point that the only evidence that should sway us is novel evidence and the more novel (or surprising) the evidence is, the more power that evidence has to change our beliefs.

5. Concluding remarks

From a very early stage, the epistemology of disagreement became chiefly concerned with how we should respond to disagreement with our peers. It did so because while responding to disagreements with non-peers seems straightforward, responding to disagreements with peers seemed far less so. What's more, some of the answers people came up with first (like the 'Equal Weight View' proposed by Adam Elga in 'Reflection and Disagreement') appeared to lead to dramatic consequences. Specifically, they seemed to make persistent disagreement on questions of prolonged study with many experts irrational. A number of philosophers noted that this result would sweep a large swathe of scholars from nearly every discipline into an irrational 'sin bin', which doesn't seem right. They then proposed alternative accounts. Those of the other side responded, and so on.

That brief sketch is, very broadly, how the epistemology of disagreement quickly became fixated on a narrow conversation about how to respond to disagreements involving peers. But the topics explored in this thesis represent a challenge to this direction. Each chapter takes an aspect of the current debate and seeks to answer one central question: does this idea help us better understand disagreement as a unique epistemic event?

It does this first for definitions of peerhood, where we discovered that most commonly-used accounts of this concept were not adequate. Then I examined the notion of 'higher order evidence' and special inference rules for disagreements, and in both cases we found that these concepts do not have a special role in disagreements. Then, when examining Lackey's arguments against Belief Independence, we found that her argument hinged on a preoccupation with epistemic peerhood, and that such a narrow application of the concept is unjustified. We have seen in each case that, after careful examination, peer disagreement may not be as unique a philosophical topic we thought it was.

It should now be clear to the reader that I think epistemology's concentration on peer disagreement is a limitation. But there is an opportunity for scholars here, because the focus on peer disagreement has meant that some other areas of disagreement have not received the attention that perhaps they deserve. I would now like to make two recommendations for areas of potential future scholarly discussion, based on some of the

limitations that I encountered when researching this paper. They are 'Opinion Pooling' and 'Model Selection'.

Opinion pooling is the problem of how best to take a collection of probabilistic judgements as inputs and reduce them to one probability judgement. It can be argued that a disagreement is a particular instance of opinion pooling where agents doing the opinion pooling also hold the opinions being aggregated. A quick survey of the opinion pooling literature brings up a lot of recognisable concepts - expert weighing, confidence, consensus, mutual information (Dietrich & List, 2013) - and these concepts are explored and discussed in the framework of Bayesian epistemology. This framework allows commentators to ask questions such as: What desiderata would such a pooling function have, and are there any proposed candidates which meet these criteria?

I think epistemology might be able to borrow two things from the discussion on opinion pooling. Firstly, it could borrow the desiderata from the opinion pooling discussion as a starting point (and perhaps a finishing point) for a disagreement response function. Following on from that it also might be able to borrow some of the conclusions from opinion pooling or at least explain why the two problems are distinct. This is the first area I would recommend as a potential focus for future scholarship.

More generally, I also think that some of the potential complexities in the epistemology of disagreement have been unexplored and this too perhaps is attributable to the concentration on peer disagreement. To illustrate, consider the following: If I ask two people who they think will win the football match this afternoon and they both pick the Bumblebees, we would normally say that they agree on this question. If I had instead asked them, what the chances are that the Bumblebees will win the match this afternoon and one of them says 60% and the other says 90% , then in fact we have uncovered a disagreement. Although their predictions about who was going to win were the same, the confidence levels underpinning their predictions were in conflict.

This disagreement was initially hidden because of the specific question we asked. This illustrates how when giving a prediction, we take our model of the world at the time and

give the most likely output of the model, yet the model we are using can remain hidden behind the prediction.

This point goes even deeper. It shows that even when we ask about the chances of x and two or more people respond with the same credence judgments, there could be potential disagreement in the models they are employing to generate those judgements. This reveals a weakness in the way that we tend to characterise disagreement in epistemology: it illustrates that whenever we approach disagreement as a clash of views on a single hypothesis or proposition, we are making a simplification and may be obscuring the underlying cause of the disagreement.

It may be that in many instances the surfacing of these underlying disagreements, were they to come to light, would impact on how someone should respond to the 'top level' disagreement. For instance, if someone told me they thought the chances of the Bumblebees winning were 99.9% then I might interpret their incredible confidence as the exuberant bias of a fan, and downgrade their epistemic status. What happens to this belief when someone tells me that they think the Bumblebees will win the match without giving me their credence? I need to account for the possibility that they are a biased fan and factor the chances of that into my assessment of their epistemic status. It is an exaggerated example, but in many disagreements there will probably be a huge number of beliefs like this that would effect the response of someone involved in the disagreement were they to come to light.

This complex belief structure reveals just how many assumptions are present when we interpret data. The resulting uncertainty illuminates the large number of explanations that are consistent with the observed data. Choosing the right model for interpreting some given information is a very consequential (and often quite a complex) task, but luckily philosophers don't need to solve this problem on their own as statisticians have spent a lot of time developing methods to address this question, often called 'Model Selection' methods.

Model Selection is the process of choosing which features we want our predictions to be optimised for. A classical example of this is Occam's Razor, which says that given two or

more theories that are consistent with our data, we should prefer the theory that is simpler. Employing Occam's Razor means that we are selecting out preferred theories (or models) both for their consistency with the data and for their parsimony.

In responding to disagreements, we apply a model of how we understand the world. But these models are dynamic and we change them in response to information as it comes in. In order to deepen our understanding of disagreement, it is necessary to reflect on which features should be preferred when we select our models for resolving disagreement. I would suggest that further research into this area could yield an enriching contribution to the epistemology of disagreement.

References:

- Christensen, D. (2007). Epistemology and disagreement, the good news, *Philosophical Review*, Vol. 116, No. 2,
- Christensen, D. (2009). Introduction: The Epistemology of Disagreement. *Episteme* 6(03): 231-232.
- Christensen D. (2009). Disagreement as Evidence: The Epistemology of Controversy. *Philosophy Compass* 4/5: 756–767
- Christensen D. (2011). Disagreement, Question-Begging, and Epistemic Self-Criticism. *Philosophers' Imprint* 11 (6):
- Dietrich, F. and List, C., 2013. Probabilistic opinion pooling generalized Part one: General agendas.
- Elga, A. (2007). Reflection and disagreement. *Noûs*, 41, 478–502.
- Elga, A. (2010). How to disagree about how to disagree. In Ted Warfield & Richard Feldman (Eds.), *Disagreement*. Oxford: Oxford University Press
- Feldman, R. (2009). 'Evidentialism, Higher-Order Evidence, and Disagreement', *Episteme*, 6(03): 294-312.
- Gelfert, A. (2011). Who is an epistemic peer? *Logos & Episteme*, 2, 507–514.
- Goldman A. (2001). Experts: Which One's Should We Trust? *Philosophy and Phenomenological Research*. 63:85-110
- Gutting, G. (1982). *Religious belief and religious scepticism*. Notre Dame: University of Notre Dame Press.
- Hájek, A., and Hartmann, S. (2010): Bayesian Epistemology, in: J. Dancy et al. (eds.), *A Companion to Epistemology*. Oxford: Blackwell, 93–106.
- Kelly, T. (2005). The epistemic significance of disagreement. In T. Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology* (Vol. 1, pp. 167–196). Oxford: Oxford University Press.
- Kelly, T. (2010). Peer disagreement and higher order evidence. In R. Feldman & T. Warfield (Eds.), *Disagreement*. Oxford: Oxford University Press.

- Kording KP, Beierholm U, Ma WJ, Quartz S, Tenenbaum JB, et al (2007) Causal Inference in Multisensory Perception. *PLoS ONE* 2(9): e943. doi:10.1371/journal.pone.0000943
- Lackey, J. (2009b). A justificationist view of disagreement's epistemic significance. In A. Haddock, A. Millar, & D. Pritchard (Eds.), *Social epistemology*. Oxford: OUP.
- Lackey, J. (2013). Disagreement and Belief Dependence, Why Numbers Matter, in Christensen and J Lackey (Eds), *The Epistemology of Disagreement: New Essays*. Oxford: Oxford University Press
- Lewis, David (1999), *Papers in Metaphysics and Epistemology*, Cambridge: Cambridge University Press.
- McGrath, S. (2008). Moral Disagreement and Moral Expertise. In Russ Shafer-Landau (Ed.) *Oxford Studies in Epistemology*. Oxford: Oxford University Press.
- Oppy, G (2010). Disagreement. *International Journal for Philosophy of Religion* 68 (1): 183-199.
- Shannon, C. E. (July 1948). A Mathematical Theory of Communication. *Bell System Technical Journal*. 27 (3): 379–423.
- Shogenji, T. A Condition for Transitivity in Probabilistic Support, *British Journal for the Philosophy of Science*, 54 (2003), pp. 613-616.
- Shogenji, T. 2007. "A Conundrum in Bayesian Epistemology of Disagreement." Unpublished manuscript.
- Simpson, R. (2012). 'Epistemic Peerhood and the Epistemology of Disagreement'. *Philosophical Studies*.
- Sosa, E. (forthcoming). The epistemology of disagreement. *The Epistemology of Disagreement: New Essays*. J. Lackey and D. Christensen, Oxford.
- Teller, P (1973), "Conditionalisation and Observation", *Synthese*, 26: 218–258.