האוניברסיטה העברית בירושלים
THE HEBREW UNIVERSITY OF JERUSALEM

THE REAL PROBLEM OF DISAGREEMENT

By

AMIR KONIGSBERG

Discussion Paper # 608 April 2012

מרכז לחקר הרציונליות

CENTER FOR THE STUDY OF RATIONALITY

Feldman Building, Givat-Ram, 91904 Jerusalem, Israel
PHONE: [972]-2-6584135 FAX: [972]-2-6513681
E-MAIL: ratio@math.huji.ac.il
URL: http://www.ratio.huji.ac.il/

The real problem of disagreement

Penultimate Version

Amir Konigsberg

The Center for the Study of Rationality, the Hebrew University, Jerusalem

1. Introduction

Recent debates surrounding the epistemic significance of peer disagreement have sought to provide responses to cases in which peers disagree about the epistemic import of a shared body of evidence. Various responses have been suggested in the literature, such as the Equal Weight View (Elga 2007), the Total Evidence View (Kelly 2009), the Common Sense View (Enoch 2010) as well as a number of other closely related approaches (Feldman 2006; Feldman 2007; D. Christensen 2007; Matheson 2009; Moss) . Typically, these address problems of the following general form:

> If persons A and B are epistemic peers – meaning roughly that it is equally probable that A and B will be correct in the domain in which they are peers – and on a particular unexpected occasion they happen to find out that they disagree about whether a particular proposition P is true given the evidence equally available to them both, and assuming that neither party has any independent reason to discount the dissenting party's conclusion, they ought to respond to this discovery … *in such and such a way*.

The responses to this problem in the literature vary, and can, I believe, be divided into three kinds: 1) the *bootstrap* response; 2) the *conciliatory* response, and 3) the *egalitarian* response. I think that some of these responses have considerable appeal. But I also think that some make sense theoretically, but are not practically sensible, and that others are simply unreasonable. All the same, I do not believe that establishing *which* of the responses presented in the literature is better, as a good deal of the recent debate has been focused on doing, actually addresses the normative problem of disagreement.

The problem of disagreement asks about the appropriate response (typically the response of a peer) upon encountering a disagreement between peers. The responses proposed in the literature offer different solutions to the problem, each of which has more or less normative appeal. Yet none of these seems to engage with what seems to be the *real* problem of disagreement. It is my aim in this paper to highlight what I think the real problem of disagreement is. It is, roughly, the problem of deciding *whether* a revisionary tactic is appropriate following the discovery of disagreement as well as deciding *which* revisionary tactic is appropriate. This, I will show, is a slippery and inevitable problem that any discussion of disagreement ought to deal with. Moreover, it is a problem that, once recognized, also impinges on the question of *which* revisionary tactic is appropriate.

The above-mentioned approaches (1-3 above) are characterized by the different tactics that they propose for dealing with disagreement. But these tactics only appear to be relevant after the truly hard work of deciding *whether* they are relevant in each actual case of disagreement has been done. And this, I believe, is a huge problem that has not been adequately recognized in the literature or has even largely been missed until now. Moreover, the epistemic significance of this problem extends far beyond debates

surrounding disagreement. It involves the subjective appreciation of evidence about the reliability of inferences from evidence in general. It is my aim here to draw attention to this problem which I believe lies at the heart of debates surrounding disagreement. It is my contention that *actual* cases of disagreement, as opposed to *possible* cases of disagreement, must deal with this inevitable situation.

The paper will proceed in two stages. In the subsequent section I will outline what I take to be the real problem of disagreement, setting forth my core argument. In section 3 I will present three approaches that characterize the solutions that have been proposed in the literature. In the course of doing so I will show why these do not address the real problem of disagreement. But I will also suggest which of the approaches in the literature is the most plausible in view of its assumed recognition of the underlying difficulty.

2 The real problem of disagreement – *setting the stage*

The crux of the matter lies in a practical paradox of sorts, which, as I shall presently show, is inevitable. It relates to judgments about evidence made from subjective standpoints such as those involved in debates about the appropriate responses to disagreement. Before I present this practical paradox I will start by clarifying what I mean by disagreement.

2.1 Disagreements between ordinary people

Disagreements as I shall refer to them are situations in which *people* disagree. More specifically, the types of disagreements I want to concentrate on, in the spirit of the recent literature, involve cases where one person finds out that another person, typically someone whose relevant epistemic capabilities are similar, holds a different opinion,

view, or belief on the same matter[1]. The fact that I will be referring to ordinary people in this context is important, and I make note of it here because it imparts that I wish to relate to the normative question of how *human reasoners* ought to respond in situations of uncertainty. This approach contrasts to another, prevalent in the economic literature, in which rational agents rather than human reasoners are the focus. The rational behavior of rational agents is typically different to that of human reasoners, and so as to set the stage for addressing the problem, the distinction must be made.

2.2 A brief note on disagreements in the economic literature

In the economic literature there has been an ongoing debate since the mid-seventies surrounding the question of disagreement[2]. The core of this debate focuses on the possibility of rational disagreement, or conflict, between rational agents. The question posed is whether it is possible that agents who are, broadly speaking, programmed to conditionalize on information in the same way, can *agree to disagree* (Aumann 1976); otherwise put, the debate considers whether it is possible for rational agents to disagree *rationally*.

The types of agents referred to in the economic literature dealing with questions concerning disagreement are *not* human agents. And the rationality that is attributed to these agents is perhaps *not* human rationality. Human rationality, in the context of disagreements, which I shall, in the spirit of the recent philosophical literature, consider here, relates to human reasoners that encounter evidence to which they may respond

[1] While opinions, views, and beliefs may suggest different meanings, each suited more than the other for a particular context; I use them here interchangeably as referring to what a person regards as true.

[2] See for instance: (Aumann 1976) and (John Geanakoplos and Heracles M. Polemarchakis 1982; Jonathan A.K. 1983; Moses and Nachum 1990; Rubinstein and Wolinsky 1990; Robin Hanson 2003; Dégremont and Roy 2009; Hansen and Cowen; Milgrom and Stokey 1982; Anon.)

imperfectly. Moreover, the problem of disagreement that I shall be concerned with is located in the wider context of human fallibility and regards disagreements as opportunities for corrective measures aimed at mitigating erroneous consequences of imperfect reasoning.

2.3 Human imperfection and its implications for practical reasoning problems

Human fallibility and the imperfection to which it gives rise can manifest itself in many ways. I shall limit myself here to discussing how it relates to resolving disagreements when these are considered in the wider context of first-person reasoning problems. In this context there are two crucial senses in which human imperfection impinges on disagreements. The first relates to the gathering of information, the second to inferring conclusions from it. People's capacities in both these areas are limited, and they commonly make mistakes. In the course of my discussion I will assume that *human* reasonsers are typically aware that they make mistakes. I will assume that they know that they are not normally capable of taking all or perhaps unlimited information into account when they are deliberating in uncertain conditions, such as those characteristic of disagreements in which it isn't clear *who* is correct. Moreover, I will also assume that typically, human reasoners are aware that when they do possess information that is relevant to their beliefs, their responses - typically their inferences from this information – are often imperfect. And by 'imperfect' in this context I have in mind, approximately, three things.

2.4 Imperfection and reasoning

Firstly, when human reasoners infer conclusions from the evidence in their possession, however limited or encompassing this evidence may be, they are not always correct in what they infer. In addition, human reasoners commonly recognize this about themselves. And what this actually means is that they recognize that their reasoning is error prone, and thus imperfect.

Secondly, while people generally know that their reasoning is error prone, they do not always recognize the occasions in which it is so. Because of this, people often think they are right when they are wrong and thus incorrect reasoning sometimes goes unrecognized.

Thirdly, because they know that they may sometimes be wrong about what they believe, and because they also know that they don't always recognize the occasions when this is so, human reasoners should not always be certain that what they believe is correct. Indeed, they ought to have some reservations about the viability of their responses to states of affairs in general and particularly when they encounter dissent from an esteemed counterpart.

2.5 When subjective credence plays against facts of the matter

In many situations of uncertainty, evidence may be more or less convincing, and this seems to play subjective credences against perhaps unknown facts of the matter. All the human reasoner has to go by is his subjective credence, which is assumed to mirror the strength of the evidence that he has. In cases such as this, an individual's subjective flaws as an evidence evaluator prescribe some type of risk mitigating strategy so that inferences made from present evidence, whose impact on credence is partly subjective, can be

weighed against some type of objective standard that is not based on the same error prone reasoning.

2.6 First-order and second-order evidence

A helpful distinction has been made in the literature, providing a convenient taxonomy for arriving at a corrective standard for the erroneous tendencies of human reasoning. The distinction is between two kinds of evidence – for simplicity let us call these first-order evidence and second-order evidence. For our present purposes the distinction should be understood as making a point about two kinds of epistemically relevant considerations that deserve epistemic appreciation and which may be taken into account in uncertain epistemic environments such as those commonly encountered by human reasoners.

In the present context *first-order evidence* refers to evidence the presence of which increases subjective credence in a particular proposition. *Second-order evidence* refers to objective evidence the existence of which makes subjective credence more or less likely to be correct. For our purposes what is entailed by first-order evidence is subjective and what is entailed by second-order evidence is objective. To flesh this out, the first-order evidence (FOE) that I encounter may be the Candlestick in the Hall which supports my belief that Colonel Mustard did it (P): <FOE |- P>. But the second-order evidence (SOE) that I possess will be my prior knowledge that in the past, when I inferred *who* was guilty on the basis of Weapon and Location alone (FOE), I was wrong 70% of the time. It is the knowledge that I have about my past performance in inferring conclusions in similar conditions (using the same variables) that provides me with second-order evidence about how likely it is that my inference <FOE |- P> is correct.

In the context of disagreements between peers, first-order evidence is the evidence that each peer encounters that consequently leads him to believe as he does in the first-place. Therefore if the disagreement is, for instance, between weather forecasters, and concerns the weather forecast for tomorrow, first-order evidence is the evidence on which each person bases his belief about tomorrow's weather. More generally, we might think of this as the type of evidence that is normally needed for a weather forecaster to make up his mind about tomorrow's weather, prior, that is, to finding out what his peer believes about tomorrow's weather. Thus first-order evidence in this context may perhaps be temperature maps (TM), atmospheric factors (AF), and other metrological features (MF) on which weather forecasters typically base their predictions. Let E denote a particular piece of evidence. In this case a prediction based on FOE will look like this: <ETM; EAF; EMF |- P>.

Second-order evidence in these contexts will typically be evidence that relates to the belief-forming circumstances in which conclusions are inferred from first-order evidence. This may for instance include considerations about how likely it is that the inferences made from first-order evidence are correct. In disagreement problems second-order evidence will typically relate to prior knowledge about a person's competence in inferring conclusions from first-order evidence (typically stated in probabilities based on prior performance).

The widely consensual position is that in reasoning problems involving first and second-order evidence, subjective credences that are based on first-order evidence ought to be balanced by probabilities derived from second-order evidence based on past performance (in the same way circumstantial indicators in Bayesian reasoning problems are weighed

against base-rate information). The general contention is that not to take into account second-order evidence, typically prior probabilities, where these are informative and thus epistemically relevant to the assessment of the viability of present evidence, is a failure of reasoning. To be more specific, it is a failure to consider objective – i.e., second-order - as well as subjective – i.e., first-order – factors, both of which are epistemically valuable. Thus to continue the example above, this would mean that if after encountering the Candlestick in the Hall I infer that Colonel Mustard did it without considering that based on second-order evidence my inference that Colonel Mustard did it is seventy percent likely to be wrong, I would be neglecting relevant and thus epistemically valuable evidence.

2.7 Theoretical and practical prescriptions

Theoretically, the prescription to consider second-order evidence in conjunction with first-order evidence appears to be unproblematic. In situations of uncertainty, second-order evidence – usually prior probabilities – ought to be weighed against subjective likelihoods derived from first-order evidence, typically in accordance with Bayes' rule[3]. But while in theory this appears unproblematic and straightforward, there seems to be a *practical* difficulty in realizing this prescription.

Because the normative prescription is that second-order evidence ought to be incorporated in reasoned deliberation in situations of uncertainty, it is subjective

[3] Bayes's rule, or theorem, is a rule for operating on numerically expressed probabilities to revise a prior probability (in other words, the base-rate) into a posterior probability after new data have been observed. According to the theorem, the posterior probability for event H1 *after* data D is observed and accounted for is: $p(H1|D) = p(H1)\ p(D|H1)/p(D)$, where $p(H1)$ is the prior probability assigned to H1 *before* D is observed.

judgment that is responsible for determining a situation *as* uncertain, and consequently for determining *whether* second-order evidence is relevant as a risk-mitigating measure for addressing this uncertainty.

Because determining whether a situation is uncertain is partly a matter that depends on how confident the person is about the first-order evidence, it can make the subjective judgment about *whether* second-order evidence is relevant dependant on the very unlawful outcome it is there to mitigate. And this is what paves the way to the *real* problem that each person that encounters second-order evidence appears to face.

2.8 Practical problems with theoretical prescriptions in disagreement problems

The practical difficulty in implementing the normative prescription to weigh first-order evidence against second-order evidence in situations of uncertainty is that from the first-person standpoint second-order evidence often has ambiguous implications. To see this, consider a hypothetical situation. Assume that I know that based on past performance I am 80% likely to be correct in my predictions. I can take this to mean that there is an 80% chance that my next prediction will be correct and a 20% chance that it won't. And having made my next prediction there appears to be no way for me to ascertain, independent of relying on my present level of conviction and the various considerations that support it, whether my present reasoning falls in the positive or negative percentiles of chance. That is to say, I have no way of knowing whether my present conviction is an instance affirming the 80% chance that I am correct or the 20% that I am not. My probability of being correct, based on past performance, is second-order evidence the inclusion of which appears to depend on my present level of confidence in the first-order

evidence. And if I am convinced about my present judgment, I may regard my present judgment as being an affirming instance of the positive likelihood of my being correct according to the second-order evidence. And because, on this interpretation, *I am* correct, there is no need to weigh my present level of confidence against second-order evidence.

This suggests that aside from the normative prescription to weigh first-order evidence against second-order evidence I appear to also have a normative obligation to weigh first-order evidence in accordance with the level of epistemic warrant that it provides. This obligation should be familiar because it appears to be standard and uncontroversial that different evidential situations warrant varying levels of confidence and that the proper incorporation of new information about these situations depends on the epistemic warrant that is provided by first-order evidence. On this reading, a person in a first-person standpoint may be faced with two, possibly conflicting, normative prescriptions. Roughly:

(1) Upon encountering a new situation respond to first-order evidence in accordance with how convincing it appears to be.

(2) Upon encountering a new situation mitigate risk under uncertainty by weighing first-order evidence against second-order evidence.

As noted, the practical problem here seems to be that there is no independent way to ascertain which prescription – (1) or (2) - applies. Moreover, this situation, in which we are asked about an individual's appropriate response once conflict with a peer is discovered, appears to be reflective of a class of epistemically ambiguous situations in which the crux seems to lie in an individual's ability to determine the appropriate

revisionary response to the situation. And because doing so is largely a matter of how confident the individual is about the evidence, the inclusion of second-order evidence seems to depend, at least in part, on the selfsame risk-prone reasoning it is there to mitigate. And this, it seems, is inevitable.

To remind the reader, the prescription to include second-order evidence in judgments concerning first-order evidence arose from the recognized imperfect responses (e.g., unpredictable conditionalization or inference) of human reasoners. Recognizing this need, second-order evidence is supposed to serve as the objective measure to mitigate the risk that may be incurred by error prone subjective judgment. Yet as we saw, as long as subjective judgment is responsible for deciding *whether* second-order evidence is relevant, no matter how 'objective' the second-order evidence may be, the problem continues.

2.9 The real problem of disagreement – a *practical* paradox

If the normative prescription is that an individual in a decision theoretic circumstance ought to decide whether or not second-order evidence is relevant to that circumstance, he must have the ability to distinguish when it is and when it isn't relevant. We assume that this is a function of how ambiguous the epistemic import from the evidence is, perhaps how weak the evidence is (Enoch 2010; Kelly 2005), or how uncertain he is about what he has inferred from the evidence. In each case it is on the basis of such considerations that the individual deems second-order evidence relevant or irrelevant. And this seems to lead to a sort of practical paradox: the judgment of relevance has no independent evidence to go by. And this means that theoretically, Bayes rule may well offer a precise

way to weigh beliefs, and philosophers writing about disagreement may well make suggestions about which responses are appropriate for peers that encounter conflicting beliefs. But tactics such as these only seem to be relevant *after* the hard work of deciding relevancy has been done, and this, as I have tried to show, is theoretically a greatly underdescribed problem that I believe has largely been missed until now[4]. It also addresses what I take to be the *real* problem of disagreement. This is my argument.

3 Responses in the literature

As noted at the outset, I believe that the proposed responses to peer disagreement in the recent philosophical literature can be divided into three kinds. Some of these seem to me to have more appeal than others, particularly in relation to the *real* problem that I outlined in section 2. In the next section I will briefly address each of the approaches in the literature. Before doing so I will sketch a pseudo-particularized example of peer disagreement on the basis of which we can assess the plausibility of each of the responses that the literature provides.

3.1 A case of peer disagreement

Jill and Jack are two equally ranked chess masters. As it happens, Jill and Jack have other things in common aside from sharing the same title and rank at chess. They have for instance both been playing chess for the same number of years and they have won the same number of games, at equally ranked tournaments, against equally classed players, using similar game strategies. Additionally, Jill and Jack also know all of this about each other.

[4] Elsewhere I discuss the epistemic significance of relevancy judgments.

On a particular occasion, Jill and Jack are each independently asked by an examiner which color has the advantage in a particular chess-board arrangement. Jill tells the examiner that she thinks that White has the advantage; Jack tells the examiner that he thinks that Black has the advantage. Then each of them is told by the examiner about what the other thinks. What should Jill and Jack do in regard to their beliefs after being given this information, assuming that is, that neither one of them has any non-question-begging reason to think that the other happens to be reasoning in sub-standard conditions – that, for instance, the other isn't drunk, dazed, tired, or anything of the sort? More specifically, should the discovery that they believe differently make either of them lose confidence in their own beliefs?

3.2 The bootstrap approach[5]

As its name suggests, the bootstrap response makes use of a person's own reasoning about the issue at hand to support his revised reasoning about the issue at hand. It suggests that because P is true it doesn't matter that John, a peer, disagrees about this, *because* P *is* true, and John is therefore wrong. This is blatantly question-begging reasoning. And it is reasoning that fails to appreciate the epistemic significance of disagreement[6]. Moreover, for disagreement to have no epistemic impact is effectively a failure to appreciate it as relevant evidence[7].

[5] Elga (Elga 2007) also discusses the problem of bootstrapping, as do Kelly (Kelly 2009) and Enoch (Enoch 2010). I do not discuss either of these views here.

[6] This is why David Enoch has fittingly called it the "*I don't care view*" ((Enoch 2010, 15); the view is attributed to Thomas Kelly (Kelly 2005). Neither Enoch nor Kelly contend that this is a plausible response to disagreement, largely because it completely ignores the epistemic significance and corrective role that other people's opinions can have on our own judgment.

[7] Matheson presents a novel argument for why evidence of disagreement is, after all, *relevant* evidence. According to Matheson, if, to continue our chess example, Jill were to ask Jack which color he thinks has the advantage on the present chess-board position, she would be justified in believing Jack on this matter,

The bootstrap response thus ignores prior knowledge concerning peerage and contends that the appropriate response to the discovery of disagreement is to act in accordance with what first-order evidence suggests[8]. With regard to the disagreement between Jill and Jack at the chess-table this would mean that Jill would take her belief that White has the advantage to support the belief that Jack is wrong *because* White has the advantage. She would thus *not* adjust her credence in White having the advantage because she has no reason to. Her disagreement with Jack appears to be epistemically insignificant, *because he is wrong*, and thus what he thinks is irrelevant.

In a wider context, the bootstrap response seems to disregard the awareness that people normally have of their own fallibility when reasoning, and the corrective role that they attribute to other people's opinions as a means for mitigating error in their own conclusions. The bootstrap response grants first-person conviction about first-order evidence a justificatory role that is normally attributed to independent and objective standards. In so doing it bootstraps the justification of a belief to first-person conviction in it.

I don't believe that anyone in the recent literature seriously considers the plausibility of this approach. It is normally referred to as a limit case expressing the tempting appeal of the kind of unwarranted reasoning that unregulated internal standards for justification can sometimes give rise to (it is, in this context, sometimes referred to as the Extra Weight

since she relies on Jack (in fact she relies on Jack on these matters as much as she relies on herself). If this is so then Jack's belief *does* provide Jill with evidence that Black has the advantage. And it is therefore evidence for Jill regardless of what she believes (Matheson 2009).

[8] In doing so it violates what Christensen has recently called "independence": "In evaluating the epistemic credentials of another's expressed belief about P, in order to determine how (or whether) to modify my own belief about P, I should do so in a way that doesn't rely on the reasoning behind my initial belief about P" (David Christensen 2011).

View. See: (Elga 2007; Kelly 2009; Enoch 2010)). All the same, what characterizes the bootstrap approach is that it only lets subjective credence in relation to first-order evidence influence its judgment. Second-order evidence has no effect. Moreover, the bootstrap approach is especially dangerous because the person using the bootstrap approach may well argue that his tactic *is* to use the outputs of first-order evidence *and* second-order evidence but that occasionally the output of first-order evidence overrides the output of second-order evidence and thus makes it irrelevant.

3.3 The conciliatory approach

The conciliatory approach is probably the most intuitively reasonable approach because it respects the corrective role that other people's beliefs can have on individual error prone reasoning. This approach recognizes both the normative requirement to respect the epistemic force of first-order evidence and the normative requirement to consider second-order evidence when things are uncertain. Accordingly, that Jack believes differently to Jill about what color has the advantage in the present alignment of the pieces on the board is epistemically significant evidence and should be incorporated into Jill's response once she finds out what Jack believes. In fact, recognizing that Jack believes differently ought to make Jill doxastically shift in the direction of Jack's belief (Matheson 2009).

The conciliatory approach contends that Jill's confidence in White being advantageously positioned should be revised following her discovery of Jack's belief. But it does not state *how* it should be revised, or more specifically, to what degree[9]. The answer to this question seems to depend on how *relevant* Jill believes Jack's belief is for her revised

[9] Matheson (Matheson 2009) clarifies and defends weak and strong conciliatory responses, from very little movements to strong movements, entailing, for instance, the suspension of judgment.

response. Or otherwise put, it depends on the measure of epistemic significance that she grants it. And here, as the reader can see, we return to the *real* problem according to which the relevance of second-order evidence depends on subjective judgments about the epistemic force of first-order evidence.

We shall return to the conciliatory approach momentarily, after examining the egalitarian approach. In the meantime let us note that the merit of the conciliatory approach appears to be that it recognizes that what an equally competent person believes *is* relevant evidence. Moreover, it prevents bootstrapped dismissals of what other people think. All the same - and perhaps here in particular lies its advantage - it leaves open the question regarding the weight that ought to be granted to prior knowledge concerning peerage. And on this matter the egalitarian response, which is a particular kind of conciliatory response, provides an answer concerning the appropriate weight that ought to be granted to second-order evidence.

3.4 The egalitarian approach

The egalitarian approach to disagreement says the following. If two people are equally likely to be correct and they unexpectedly discover that they hold different beliefs about what *is* correct then (assuming that they have nothing that is independent of their reasoning about first-order evidence to back this up) they ought to recede their confidence in their own belief being correct to the conditional probability that would be granted to their being correct in considering, prior to the actual disagreement, the appropriate response were such a situation to occur[10]. And because they are equally likely

[10] This is largely based on Elga's formulation for the Equal Weight View (Elga 2007): "Upon finding out that an advisor disagrees, your probability that you are right should equal your prior conditional probability

to be correct, the appropriate conditional probability, were such a situation to occur, would be 0.5. Why 0.5? Because from the theoretical vantage point the probability that they would be correct if such a situation would occur is equal. Otherwise put, if two peers disagree and neither has any reason aside from the other's opinion to think that the other is wrong, such as their possession of more or better evidence or their superior conditions for inferring from the evidence in the present circumstance, then the epistemic weight of first-order evidence ought to be ignored, and the revised credence ought to be based on second-order evidence alone.

The egalitarian approach suggests what Sarah Moss has called a "perfect compromise" between agents that possess differing credences to propositions. On this suggestion, if A assigns credence C_1 to P, and B assigns credence C_2 to P, then a 'perfect compromise' would be for A and B to assign $(C_1+ C_2)/2$ credence to P (Moss forthcoming). Note that the egalitarian approach amasses its conclusions from theoretical thinking about a *possible* disagreement occurring. It contends that if two people know that based on their prior performance and capabilities they share an equal probability of being correct on some hypothetical occasion, and they also consider (from this same hypothetical perspective) the probability of their each being correct on the occasion of such a hypothetical disagreement occurring, then their *actual* response ought to be equal to their hypothetical response, that is, treating each belief as equally probable.

The egalitarian approach doesn't take into account that in an *actual* case of disagreement each of the parties involved has reasons for believing as they do and consequently also

---

that you would be right. Prior to what? Prior to your thinking through the disputed issue, and finding out what the advisor thinks of it. Conditional on what? On whatever you have learned about the circumstances of the disagreement" (Elga 2007).

reasons for dismissing the purely hypothetical suggestion. The actual disagreement as opposed to the theoretical disagreement appears to possess an epistemic factor that the theoretical consideration disregards. In the actual disagreement one person may have reason to think that they are correct, and as a consequence of this to think that the other person is incorrect[11]. From the theoretical vantage point this kind of playing of subjective credence against second-order evidence is unwarranted. From the actual perspective of a person in the midst of disagreement the justificatory weight of subjective credence is not only plausible but also appears to be a normative requirement – namely a requirement to respond to the evidence in a manner that is sanctioned by *that* evidence. To dismiss first-order evidence altogether, as egalitarian positions such as the Equal Weight View (Elga 2007) and *splitting the difference* (D. Christensen 2007) seek to do, leads to a sort of real-world skepticism (Feldman 2006, 415), according to which the level of credence attributed to many of our commonly held beliefs ought to be reduced, which is implausible.

Adam Elga's (Elga 2007) and David Christensen's (D. Christensen 2007) views are two examples of what I take to be the egalitarian approach to peer disagreement. On both of these accounts, if disagreement is apparent, each peer ought to revise their confidence in alignment with what second-order evidence dictates. Thus upon discovering that Jack believes that Black has the advantage, Jill ought to recede the confidence that she has in her belief to what it would have been conditional on the hypothetical possibility of such a disagreement occurring. She thus ought to give 'equal weight' to her belief and to Jack's

[11] And in the actual disagreement this might be explained in a number of ways. The dissenting person, in view of his divergence of opinion, may be thought to have slipped in performance - made a mistake that is, a performance-error, perhaps misapplying the proper rules of inference. Alternatively, the divergence itself may be regarded as a reason, or perhaps even a proof, that that person ought to be demoted from the level of peer.

in revising her belief (Elga 2007). Otherwise put, she ought to 'split the difference' between what she believes and what Jack believes (D. Christensen 2007).

The merit of these egalitarian suggestions is that they obey prior probabilities in the absence of any non-question-begging circumstantial indicators. Their limitation is that they eradicate the epistemic weight of the first-order considerations on the basis of which Jill came to believe as she did in the first place. Otherwise put, these approaches contend that if there is to be compromise the only compromise there can be is perfect compromise, while actual disagreements suggest that it isn't at all clear that this is so.

There is a serious danger in the egalitarian approach. It is the danger of rigid reasoning. If a person does not honors the prior probabilities entailed by second-order evidence but instead defers to them completely, his decisions and behavior will be more rigid. New experiences will be classified on the basis of previously established probabilities and information will be absorbed less for its intrinsic value and more according to whether it meets the terms of a bet.

3.5 Back to the conciliatory approach

Once we recognize that actual disagreements involve reasoning on the basis of everything that is epistemically available to us (Enoch 2010, 38), including subjective contentions regarding the conclusiveness of first-order evidence, the dangers of bootstrapping seem to be inevitable. All the same, to completely overrule first-order evidence on the basis of second-order evidence, as the egalitarian approach suggests, seems wrong. In this sense the dangers of bootstrapping are dangers that we will perhaps have to live with (Enoch 2010). But without preserving an individual's normative obligation to deal with these

dangers we would seem to have to give up too much of what good reasoning amounts to. Indeed, to eradicate the possibility of error altogether by dismissing any possibility of inaccuracy would seem to compromise our ability to encounter and conditionalize on new evidence.

In sum I think that the conciliatory approach provides the most plausible tactic for responding to disagreement. It is a common-sense approach that accepts the corrective role that disagreement can have and contends that ignoring what other people think is, sometimes, irrational. The fact that what someone else believes is relevant and epistemically valuable evidence appears to be an important factor in any person's long-term strategy for avoiding error. That the conciliatory approach doesn't specify a rule for how to balance first-order evidence with second-order evidence appears to leave much of the hard work to the individual's sense of judgment. In doing so it is perhaps less explicit in its rules for revision but it is, nonetheless, more attuned to the complexity involved.

4 So what, after all, can we do?

That the deliberating subject is responsible for mitigating his own imperfect responses by second-order evidence is inevitable. As we saw, this can be dangerous because precautious strategies may be based on circular reasoning. All the same, to dismiss subjective judgment altogether upon encountering disagreement is implausible because it will entail that (I) many beliefs about which we are perfectly confident will have to be dismissed and (II) it will undermine a subject's ability to attain new information whenever conflict is encountered. That being said, the appropriate response ought to honor both the normative requirement to consider second-order evidence in uncertain

situations and to respect the first-order evidence on the basis of the epistemic force it provides. How these two requirements should be considered in an actual case of belief revision – how much weight ought to be granted to each - is something that must be left to the subjective judgment of the subject that is aware of these two requirements. And this, as I said at the outset, is inevitable.

References:

Anon. Agreeing To Disagree: A Survey. http://ideas.repec.org/p/fth/caldec/97-18.html.

Aumann, Robert J. 1976. "Agreeing to Disagree." *The Annals of Statistics* 4 (6) (November 1): 1236-1239.

Christensen, D. 2007. "Epistemology of Disagreement: The Good News." *Philosophical Review* 116 (2) (April): 187-217. doi:10.1215/00318108-2006-035.

Christensen, David. 2011. "Disagreement, Question-Begging, and Epistemic Self-Criticism." *Philosophers' Imprint* 11 (6): unknown.

Dégremont, Cédric, and Olivier Roy. 2009. Agreement Theorems in Dynamic-Epistemic Logic. In *Logic, Rationality, and Interaction*, ed. Xiangdong He, John Horty, and Eric Pacuit, 5834:105-118. Berlin, Heidelberg: Springer Berlin Heidelberg. http://www.springerlink.com/content/g037338q80122l85/.

Elga, Adam. 2007. "Reflection and disagreement." *Noûs* 41 (3): 478–502.

Enoch, David. 2010. "Not Just a Truthometer: Taking Oneself Seriously (but not Too Seriously) in Cases of Peer Disagreement." *Mind* 119 (476) (October 1): 953 - 997. doi:10.1093/mind/fzq070.

Feldman, Richard. 2006. Epistemological puzzles about disagreement,. In *Epistemology Futures,*, ed. Stephen Hetherington, 216-326,. Oxford University Press,.

———. 2007. Reasonable religious disagreements,. In *Philosophers Without Gods: Meditations on Atheism and the Secular,*, ed. Louise Antony, 194-214,. Oup,.

Hansen, Robin, and Tyler Cowen. "Are disagreements honest?" *Journal of Economic Methodology* to appear. http://hanson.gmu.edu/deceive.pdf.

John Geanakoplos, and Heracles M. Polemarchakis. 1982. *We Can't Disagree Forever*. Cowles Foundation for Research in Economics, Yale University, July. http://ideas.repec.org/p/cwl/cwldpp/639.html.

Jonathan A.K., Cave. 1983. "Learning to agree." *Economics Letters* 12 (2): 147-152. doi:10.1016/0165-1765(83)90126-X.

Kelly, Thomas. 2005. The epistemic significance of disagreement. In *Oxford Studies in Epistemology, Volume 1*, ed. John Hawthorne and Tamar Gendler Szabo, 167-196. Oup.

———. 2009. Peer disagreement and higher order evidence,. In *Disagreement,*, ed. Ted Warfield and Richard Feldman. Oup,.

Loewenstein, George. 1992. *Choice over time*. Russell Sage Foundation.

Matheson, Jonathan. 2009. “Conciliatory Views of Disagreement and Higher-Order Evidence.” *Episteme* 6 (October): 269-279. doi:10.3366/E1742360009000707.

Milgrom, Paul, and Nancy Stokey. 1982. “Information, trade and common knowledge.” *Journal of Economic Theory* 26 (1) (February): 17-27.

Moses, Yoram, and Gal Nachum. 1990. Agreeing to disagree after all. In *Proceedings of the 3rd conference on Theoretical aspects of reasoning about knowledge*, 151–168. TARK ’90. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc. http://dl.acm.org/citation.cfm?id=1027014.1027037.

Moss, Sarah. forthcoming. “Scoring Rules and Epistemic Compromise.” *Mind*

Robin Hanson. 2003. “For Bayesian Wannabes, Are Disagreements Not About Information?” *Theory and Decision* 54 (2). Theory and Decision: 105-123.

Rubinstein, Ariel, and Asher Wolinsky. 1990. “On the logic of ‘agreeing to disagree’ type results.” *Journal of Economic Theory* 51 (1). Journal of Economic Theory: 184-193.