



AARHUS UNIVERSITY



Coversheet

This is the accepted manuscript (post-print version) of the article.

Contentwise, the accepted manuscript version is identical to the final published version, but there may be differences in typography and layout.

How to cite this publication

Please cite the final published version:

Rasmussen, M. S., Steglich-Petersen, A., & Bjerring, J. C. K. (2018). A Higher-Order Approach to Disagreement. *Episteme*, 15(1), 80-100. DOI: 10.1017/epi.2016.43

Publication metadata

Title:	A Higher-Order Approach to Disagreement
Author(s):	Mattias Skipper Rasmussen, Asbjørn Steglich-Petersen & Jens Christian Bjerring
Journal:	Episteme
DOI/Link:	10.1017/epi.2016.43
Document version:	Accepted manuscript (post-print)

General Rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

If the document is published under a Creative Commons license, this applies instead of the general rights.

A Higher-Order Approach to Disagreement

Mattias Skipper Rasmussen*
Asbjørn Steglich-Petersen
Jens Christian Bjerring

Penultimate draft
Forthcoming in *Episteme*

Abstract

While many philosophers have agreed that evidence of disagreement is a kind of higher-order evidence, this has not yet resulted in formally precise higher-order approaches to the problem of disagreement. In this paper, we outline a simple formal framework for determining the epistemic significance of a body of higher-order evidence, and use this framework to motivate a novel interpretation of the popular “equal weight view” of peer disagreement—we call it the Variably Equal Weight View (VEW). We show that VEW differs from the standard Split the Difference (SD) interpretation of the equal weight view in almost all cases of peer disagreement, and use our formal framework to explain why SD has seemed attractive but is in fact misguided. A desirable feature of VEW, we argue, is that it gives rise to plausible instances of *synergy*—an effect whereby the parties to a disagreement should become more (or less) confident in the disputed proposition than any of them were prior to disagreement. Lastly, we show how VEW may be generalized to cases of non-peer disagreement.

Keywords Epistemology of disagreement · Peer disagreement · Higher-order evidence · Higher-order defeat · Equal weight view · Conciliationism

*Address for correspondence: Mattias Skipper Rasmussen, Department of Culture and Society, Aarhus University, e-mail: filmsr@cas.au.dk

1 Introduction

When you find yourself disagreeing with an epistemic peer, how (if at all) should you revise your credence in the disputed proposition? This question has generated an extensive literature over the past decade. Following custom, we can broadly divide different views of disagreement into a “steadfast” camp and a “conciliatory” camp. Roughly, according to steadfast views, one is typically permitted to stick to one’s guns and maintain one’s credence in the face of peer disagreement. By contrast, conciliatory views maintain that one is typically required to revise one’s credence in the direction of one’s peer.¹

For all their differences, there is a broad consensus among contenders in the disagreement debate that evidence of disagreement is a kind of *higher-order evidence*—roughly, evidence about the existence or bearing of ordinary first-order evidence.² There is, however, little consensus about how to understand the normative impact of higher-order evidence. On the one hand, David Christensen has argued that higher-order evidence can require an agent to “put aside” or “bracket” her ordinary first-order evidence, thereby preventing the agent from giving the original first-order evidence its due (Christensen 2010, p. 195). In particular, on Christensen’s view, the higher-order evidence in cases of peer disagreement serves to cancel or undermine the normative impact of the first-order evidence. This, in turn, naturally aligns with the conciliatory dictum that parties to a peer disagreement should revise their opinions in a uniform way, regardless of who judged the first-order evidence correctly to begin with. On the other hand, Thomas Kelly has argued that the higher-order evidence in cases of peer disagreement does *not* generally “swamp” or undermine the normative impact of the first-order evidence (Kelly 2010, p. 192). This view, in turn, aligns with the steadfast dictum that disagreeing peers should revise their opinions in a way that properly reflects

¹See Kelly (2005; 2010), Van Inwagen (2010), and Titelbaum (2015) for some broadly steadfast views of disagreement. For defenses of conciliationism, see Christensen (2007a), Elga (2007), Feldman (2006), Matheson (2009; 2015), among others. Lackey (2008a; 2008b) and Worsnip (2014) both present “moderate” views of disagreement that purport to navigate between the steadfast and the conciliatory camps.

²For authors who discuss disagreement under the heading of higher-order evidence, see Christensen (2010), Feldman (2009), Kelly (2005; 2010), Matheson (2009), among others.

who judged the original first-order evidence correctly to begin with. As such, the divide between steadfast and conciliatory views of disagreement seems to stem, at least partly, from a more general divide over how to understand the epistemic significance of higher-order evidence. It is therefore natural to expect that general considerations about higher-order evidence may help us to better understand and adjudicate between different views of the epistemic significance of disagreement.

In this paper, we outline (§2) a simple formal framework for determining the epistemic significance of a body of higher-order evidence, and use this framework (§3) to motivate a novel interpretation of the popular “equal weight view” of disagreement—we call it the Variably Equal Weight View (VEW). According to VEW, disagreeing peers should always be given equal weight, but the amount of weight they should be accorded varies depending on how *reliable* they are: roughly, reliable peers should be given equal *high* weight, and unreliable peers should be given equal *low* weight. We show (§4) that VEW differs from the standard Split the Difference (SD) interpretation of the equal weight view in almost all cases of disagreement, and use our formal framework to explain why SD has seemed attractive but is in fact misguided. A desirable feature of VEW, we argue (§5), is that it gives rise plausible instances of what Easwaran et al. (forthcoming) call *synergy*—an effect whereby the parties to a disagreement should become more (or less) confident in the disputed proposition than any of them were prior to the disagreement. Lastly, we generalize VEW to cases of non-peer disagreement (§6).

2 Higher-Order Evidence

Let us begin by characterizing higher-order evidence more precisely. As of yet, there is no single, generally agreed upon definition of higher-order evidence in the literature, but higher-order evidence has been variously characterized as “evidence of my own rational failure” (Christensen 2010, p. 185), “[e]vidence about the existence, merits, or significance of a body of evidence” (Feldman 2009, p. 304), or “evidence that suggests that the bearing of the

original, first order evidence is something other than what one initially took it to be” (Kelly 2010, p. 200).³ To illustrate, consider a paradigmatic higher-order case adapted from Christensen (2010, p. 187):

Performance-diminishing drug: Suppose I carefully prove a logical theorem T . Since I know myself to be a competent logician, I become confident that T is in fact a theorem. However, my doctor now calls to tell me that she has given me a performance-diminishing drug that subtly but significantly impairs my reasoning abilities. In light of this new information, I become less confident that my proof is sound and that T is in fact a theorem.

Here, the doctor’s testimony seems to rationalize a change of belief in T , not by speaking directly for or against T , but by indicating that my reasons for being confident in T are weaker than I initially took them to be. So the doctor’s testimony is higher-order evidence *against* T . Conversely, if we had replaced the performance-diminishing drug by a performance-enhancing drug, the doctor’s testimony would have been higher-order evidence *for* T .

More generally, we will say that a body of evidence is *first-order* iff it bears on what we will call a *first-order hypothesis*: an ordinary proposition such as “the sun is shining” or “ $2 + 3 = 5$.” Similarly, we will say that a body of evidence is *higher-order* iff it bears on what we will call a *higher-order hypothesis*: a proposition about an evidential relation between a first-order hypothesis and a body of first-order evidence. We can think of higher-order hypotheses as propositions of the form “this-or-that evidence bears on this-or-that hypothesis in this-or-that way.” For example, the hypothesis that “my calculations do not support that T is a theorem” is higher-order in virtue of being about an evidential relation between a body of first-order evidence (my calculations) and a first-order hypothesis (“ T is a theorem”).

Formally, we will represent higher-order hypotheses using a standard probability function P which, given a hypothesis and a body of evidence, yields a sharp probability of the hypothesis given the evidence. For example,

³See Lasonen-Aarnio (2014; 2015), Matheson (2009), Schoenfield (2015), and Whiting (2016) for similar characterizations of higher-order evidence.

we will write $P(H|E) = .57$ to say that the probability of hypothesis H given evidence E is .57.⁴ We will assume that rational credences conform uniquely to P —that is, a rational agent with total evidence E bearing on H has a credence in H that is identical to $P(H|E)$.⁵ This assumption, together with the assumption that evidential probabilities are sharp, implies what Feldman (2007) has labelled “the uniqueness thesis:” for any body of evidence bearing on some hypothesis, there is a unique credence that is rational to have in that hypothesis given that evidence. The equal weight view is often thought to be tightly connected to the uniqueness thesis. For instance, Feldman (2007) argues that the uniqueness thesis commits one to the equal weight view, and Kelly (2010) argues, conversely, that the equal weight view is implausible (partly) because it presupposes the uniqueness thesis. Yet, it remains contentious how (if at all) the equal weight view is connected to the uniqueness thesis, and we will not pursue this question here.⁶

Another central distinction will be that between first-order and higher-order *probabilities*. If H is a first-order hypothesis and H_H is a higher-order hypothesis, then $P(H|\cdot)$ is a first-order probability, and $P(H_H|\cdot)$ is a higher-order probability. For example, we write $P(P(H|\cdot) = .57 | \cdot) = .3$ to say that there is a higher-order probability of .3 that the first-order probability of H is .57. For ease of notation, we will henceforth use a higher-order

⁴In fact, we doubt that evidential probabilities generally obey the standard probability axioms, at least in cases of misleading higher-order evidence. For example, in *Performance-diminishing drug*, even if my proof of T is in fact sound, it seems rational for me to lose confidence in T upon receiving evidence that I was cognitively impaired while constructing the proof. This seems to suggest that the evidential probability of tautologies can be lower than 1. If so, it follows that evidential probabilities can violate the standard probability axioms. Yet, as it is well beyond the scope of this paper to assess the costs and benefits of giving up the standard Bayesian account of evidential probabilities, we will follow custom and use a probability function to reason about evidential probabilities.

⁵By assuming that rational credences conform to P , we should not be taken to say (implausibly) that what it *means* to have a credence of p in a hypothesis H is to believe that the probability of H given one’s evidence is p . There is no assumption that credences can be reduced to—or must be accompanied by—binary beliefs with a probabilistic content. Thanks to an anonymous reviewer for pressing us to clarify this point.

⁶For recent discussions of the uniqueness thesis, see Ballentyne & Coffman (2011; 2012), Christensen (2014), Elga (2010), Greco & Hedden (forthcoming), Levinstein (2013, ch. 4), Schoenfield (2013), and White (2005).

probability function defined as

$$Q(p) =_{df} P(P(H|\cdot) = p | \cdot),$$

where H is a given first-order hypothesis, and p is a probability variable. Given this, we can simply write $Q(.57) = .3$ to say that there is a higher-order probability of .3 that the first-order probability of H is .57. As usual, first-order and higher-order probabilities may or may not be conditional on evidence.

If we know the higher-order probability distribution $Q(p)$ for a first-order hypothesis H , the Law of Total Probability (LoTP) offers a simple way of determining the first-order probability of H :

$$P(H|\cdot) = \sum_{i=1}^n p_i Q(p_i) \tag{LoTP}$$

This formulation of LoTP presupposes that p is a discrete probability variable that can take on n different values p_1, \dots, p_n . If p is continuous, the finite sum above will instead be an integral over the unit interval $[0, 1]$, and $Q(p)$ will be a density function. Here is a simple example of how to use LoTP:

Two kinds of dice: You are about to roll a random die from a bowl of dice, and you have conclusive evidence (E_H) that half of the dice are three-sided, while the other half are six-sided. Before picking a die, you consider the likelihood that (H) you will roll a six. In light of your evidence, you are rationally certain that there is a .5 chance that the die is three-sided, and similarly a .5 chance that it is six-sided. And while there is no chance of rolling a six if the die is three-sided, there is a one sixth chance of doing so if the die is six-sided. So $Q(0) = Q(\frac{1}{6}) = .5$. By plugging these higher-order probabilities into LoTP, we can calculate the first-order probability of H given E_H as follows: $P(H|E_H) = 0 \cdot .5 + \frac{1}{6} \cdot .5 \approx 8\%$.

The central benefit of this formal setup is that it allows us to capture how a body of higher-order evidence can make a difference to the evidential probability of a first-order hypothesis H , not by speaking directly for

or against H , but by influencing the shape of the higher-order probability distribution for H . Obviously, the question remains what the higher-order probability distribution $Q(p)$ will look like given different bodies of evidence, and, in particular, how we should determine $Q(p)$ in cases where an agent possesses a mixed body of first-order and higher-order evidence. But if we can answer this question, LoTP can help us calculate the resulting first-order probability.

In the next section, we will use the core intuition behind the equal weight view to specify the higher-order probability distribution $Q(p)$ in cases of peer disagreement. We choose to focus our attention on the equal weight view partly because of its prominence in the literature, and partly because it remains the best game in town to our mind. Yet, we are optimistic that our higher-order approach may be used to interpret other positions in the disagreement debate as well.

3 The Variably Equal Weight View

Consider a paradigmatic case of peer disagreement adapted from Levinstein (2015, p. 1):

Weather forecasters: Each morning, two equally experienced weather forecasters assess the same body of meteorological data and make private predictions of whether it will rain later that day. This morning, they learn that their predictions conflict: one has predicted that it is going to rain, while the other has predicted that it is not going to rain. How (if at all) should they revise their weather forecasts?

According to the equal weight view, in revising their weather forecasts, the disagreeing weather forecasters should place *equal weight* on each other's predictions. The driving intuition is that, absent reasons to think that I'm in a better epistemic position than you with respect to some proposition, it seems arbitrary and question-begging of me to place more weight on my own

opinion just because it's mine.⁷

Compelling as it is, the equal weight dictum, in its bare form, underdetermines how disagreeing peers should revise their opinions. For the mere requirement that disagreeing peers should place equal weight on each other's opinions leaves open just how *much* weight they should so place. For example, the weather forecasters from above can strictly speaking satisfy the equal weight dictum by placing *no* weight at each other's predictions whatsoever. Yet, this would seem to go against the spirit of the equal weight dictum. On a more plausible interpretation, the weather forecasters should place an amount of weight on each other's predictions that properly reflects how competent or *reliable* they are at judging the meteorological data: if they are reliable at judging the data, they should place equal *high* weight on each other's predictions, and if they are *unreliable* at judging the data, they should place equal *low* weight on each other's predictions.

In light of these considerations, we propose the following qualification of the equal weight dictum:

The variably equal weight dictum: Disagreeing peers should always place equal weight on each other's opinions, but the amount of weight they should accord each other depends on how reliable they are at judging the shared first-order evidence.

In the remainder of this section, we will use the formal tools introduced in §2 to make the variably equal weight dictum precise. As a first step, we need a general characterization of peer disagreement cases that specifies how reliable the disagreeing peers are at judging the available evidence. To this end, we will use the following template:

Peer disagreement: Two agents *a* and *b* share their total body of first-order evidence *E* bearing on a first-order hypothesis *H*. The agents are rationally certain that they both have a reliability of *r*—that is, they are both certain that there is a probability of *r* that *a* will judge *E* correctly, and

⁷For detailed defenses of the equal weight view, see Christensen (2007a), Elga (2007), Matheson (2015), among others.

similarly for b . The agents form independent opinions about the bearing of E on H : a judges that $P(H|E) = p_a$, and b judges that $P(H|E) = p_b$. Which credence should the agents adopt in H upon disagreement?

Let us clarify a few points about the notion of reliability used here. We represent an agent’s reliability by a real number $r \in [0, 1]$ that denotes the probability, prior to judging a body of evidence, that the agent will judge the evidence correctly (i.e. make a judgment identical to the true evidential probability). For example, if an agent with a reliability of .3 judges that $P(H|E) = .9$, then there is a probability of .3 that, in fact, $P(H|E) = .9$, and, accordingly, a probability of .7 that $P(H|E) \neq .9$.⁸ This notion of reliability is intentionally idealized. Presumably, agents like you and I rarely (if ever) judge evidence in this highly precise way. Later, in §6, we introduce a more sophisticated model of reliability which is roughly designed to represent the probability that an agent’s judgment is at least “in the right ballpark.” But for now, we stick to the idealized representation to keep the formalism simple.

Our notion of reliability is “evidence-based” in the sense that it captures an agent’s ability to judge the evidence at hand. As such, it differs importantly from “truth-based” notions of reliability that capture how well an agent’s credences track the *truth*. Truth-based notions of reliability are useful for a number of theoretical purposes. For example, Easwaran et al. (forthcoming) combine a truth-based notion of reliability with Bayes’ theorem to derive a simple heuristic for how to adjust one’s credence in response to learning the credences of others (we return to Easwaran et al. in §5 below). Yet, given the usual assumption in the peer disagreement literature that peers are equally good at judging the available evidence, we take it that our evidence-based notion of reliability sits well with the standard usage in the disagreement debate.⁹ This also seems like the more relevant notion of reliability given the usual stipulation that epistemic peers share a body of

⁸Note that the fact that there is a probability of .3 that $P(H|E) = .9$ does not imply that there is a probability of .3 that, for example, $P(H|E) \neq .2$ —this probability will be considerably higher. Thanks to an anonymous reviewer for drawing our attention to this point.

⁹For authors in the peer disagreement literature who operate with an evidence-based notion of reliability, see Christensen (2007a), Kelly (2010), Lackey (2008a), and others.

evidence which they judge independently of each other.

To reach a verdict in the peer disagreement case, the task is then to determine the probability of the first-order hypothesis H given the peers' total evidence E_T consisting of the first-order evidence E in combination with the higher-order evidence given by the peers' judgments. We know from §2 that if we can specify the higher-order probability distribution $Q(p)$ for H in *Peer disagreement*, then we can use LoTP to determine $P(H|E_T)$.¹⁰ To specify $Q(p)$, note first that two peers a and b may have come to disagree in either of three ways: either (i) a judged the evidence correctly and b didn't, or (ii) b judged the evidence correctly and a didn't, or (iii) both a and b misjudged the evidence. The only scenario that is ruled out by the disagreement between a and b is the one where *both* peers judge the evidence correctly. We can think of the likelihood of scenario (i) as representing the higher-order probability $Q(p_a)$ that the first-order probability $P(H|E)$ is identical to a 's judgment p_a . Similarly, the likelihood of scenario (ii) represents the higher-order probability $Q(p_b)$ that the first-order probability $P(H|E)$ is identical to p_b . Finally, the likelihood of scenario (iii) represents the sum of the higher-order probabilities $Q(p_i)$, for $p_a \neq p_i \neq p_b$. So to specify the higher-order probability distribution $Q(p)$, we need to determine the likelihood of each of the scenarios (i)-(iii) from the point of view of the peers upon disagreement.

Given the supposition that a and b have a shared reliability of r , there is a probability of $1 - r$ that a misjudges the evidence, and similarly for b . So there is a probability of $(1 - r)^2$ that both a and b have misjudged the evidence, and, accordingly, a probability of $1 - (1 - r)^2$ that one of the peers has judged the evidence correctly. That is, the likelihood that scenario (iii) obtains is $(1 - r)^2$, and the likelihood that either scenario (i) or (ii) obtains is $1 - (1 - r)^2$. Moreover, given the dictum that disagreeing peers should be given equal weight regardless of whom (if any) of them actually judged

¹⁰Note that Bayes' theorem does not provide an easy way of determining $P(H|E_T)$ in the current setting. For since we operate with an evidence-based notion of reliability, we do not have a direct link between an agent's credence in H and the truth-value of H . Rather, we merely have a link between the agent's credence in H and the evidential probability of H given the agent's (possibly misleading) evidence. As such, an agent's reliability is not useful for determining the inverse probability $P(E_T|H)$ which is needed to apply Bayes' theorem.

the evidence correctly, we should divide the probability $1 - (1 - r)^2$ evenly on scenarios (i) and (ii). As such, we get the following higher-order probabilities:

$$Q(p_a) = Q(p_b) = \frac{1 - (1 - r)^2}{2}. \quad (1)$$

This equation tells us that the higher-order probability of the higher-order hypothesis $P(H|E) = p_a$ is $.5(1 - (1 - r)^2)$, and similarly for $P(H|E) = p_b$. As we should expect, the value of $Q(p_a)$ and $Q(p_b)$ increases when r increases—that is, the more reliable an agent is, the more likely the agent is to judge the first-order evidence correctly.

Consider, next, the remaining higher-order probabilities $Q(p)$, for $p_a \neq p \neq p_b$. Given that the probability variable p is discrete and can take on n different values, there are a total of $n - 2$ such remaining higher-order probabilities. Thus, the likelihood of scenario (iii)—namely $(1 - r)^2$ —should be distributed on $n - 2$ different higher-order probabilities. What should this distribution look like? Given our idealized model of reliability, the peers' judgments provide us with information about the value of $Q(p_a)$ and $Q(p_b)$, but not about the value of $Q(p)$, for $p_a \neq p \neq p_b$.¹¹ So we need to assign prior values to $Q(p)$, for $p_a \neq p \neq p_b$. One simple and intuitive way of doing this is by using the well-known Principle of Indifference (PoI) which says that, absent reasons to the contrary, one should assign equal probability to competing hypotheses.¹² By invoking PoI, we can distribute the probability $(1 - r)^2$ evenly on each $Q(p)$, for $p_a \neq p \neq p_b$:

$$Q(p) = \frac{(1 - r)^2}{n - 2} \quad \text{for } p_a \neq p \neq p_b. \quad (2)$$

As we should expect, the value of $Q(p)$, for $p_a \neq p \neq p_b$, decreases when r increases: the more reliable an agent is, the less likely the agent is to misjudge the first-order evidence.

¹¹As already mentioned, this way of thinking about reliability is intentionally simplified. More realistically, if an agent is highly reliable, the value of $Q(p)$ will presumably be higher when p is close to the agent's judgment than when p is far from the agent's judgment. In §6, we refine our model of reliability to capture this idea.

¹²See Keynes (1921) and Jaynes (1973; 2003) for some classic discussions of PoI. See Pettigrew (2016) for a recent defense of PoI.

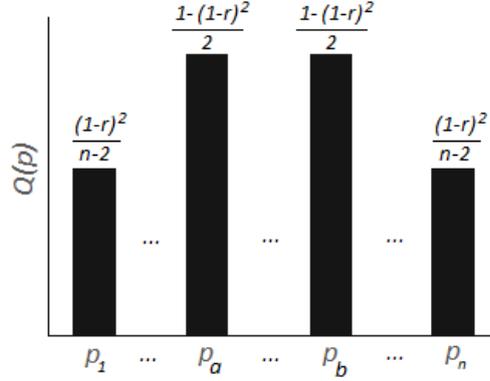


Figure 1: The higher-order probability distribution in *Peer disagreement*.

Taken together, (1) and (2) constitute the higher-order probability distribution in *Peer disagreement* (see figure 1 for an illustration). If we plug this distribution into LoTP, we get the following first-order probability of H given the total evidence E_T :

$$P(H|E_T) = \frac{1 - (1-r)^2}{2} \cdot (p_a + p_b) + (1-r)^2 \cdot \sum_{i=1, a \neq i \neq b}^n \frac{p_i}{n-2}. \quad (3)$$

This formula dictates which credence a and b should adopt upon disagreement given that p is discrete. But since credences are typically allowed to take on any value between 0 and 1, we are really interested in the case where p is continuous in the interval $[0, 1]$. To this end, we let n approach infinity to arrive at our formal interpretation of the variably equal weight dictum:

The Variably Equal Weight View (VEW): The parties to a peer disagreement should adopt a credence identical to the following evidential probability:

$$P(H|E_T) = \frac{1 - (1-r)^2}{2} \cdot (p_a + p_b) + \frac{(1-r)^2}{2},$$

where H is the disputed first-order hypothesis, r is the peers' shared reliability, and E_T is the peers' total evidence consisting of the first-order evidence E and the higher-order evidence given by the peers' judgments p_a and p_b of the evidential probability of H given E .

We can think of the different components of the VEW equation in terms of the scenarios (i)-(iii): the first summand (the one that features p_a and p_b) corresponds to the scenarios (i) and (ii) where one of the peers has judged the evidence correctly, and the second summand (the one that doesn't feature p_a and p_b) corresponds to scenario (iii) where both peers have misjudged the evidence. The relative impact or weight of these two summands on the value of $P(H|E_T)$ depends on the relative likelihood of the scenarios (i)-(iii). These likelihoods, in turn, depend on the peers' shared reliability: high values of r leads to a high likelihood of (i) and (ii) and a low likelihood of (iii), and vice versa. Accordingly, the impact that the peers' judgments p_a and p_b have on the value of $P(H|E_T)$ depends on the value of r . We may naturally depict this dependency in terms of the derivative of $P(H|E_T)$ with respect to the sum of p_a and p_b (see figure 2):

$$\frac{\partial P(H|E_T)}{\partial(p_a + p_b)} = \frac{1 - (1 - r)^2}{2}.$$

Intuitively, $\partial P(H|E_T)/\partial(p_a + p_b)$ represents the “impact factor” of p_a and p_b on $P(H|E_T)$. This impact factor is solely a function of the peers' shared reliability: high values of r result in a high impact factor of p_a and p_b on $P(H|E_T)$, and low values of r result in a low such impact factor. This reflects the spirit behind the variably equal weight dictum: reliable disagreeing peers should give each other equal high weight, while unreliable disagreeing peers should give each other equal low weight.

In the extreme case where $r = 0$, the peers should place no weight on each other's judgments: a change in the value of p_a and p_b is not going to make any difference to the value of $P(H|E_T)$. By contrast, when $r = 1$, the peers should place maximum weight on each other's judgments: if the average value of p_a and p_b changes by a given amount, the value of $P(H|E_T)$ is going to change by the same amount. For example, if the average value of p_a and p_b increases by .3, the value of $P(H|E_T)$ also increases by .3, and if the average value of p_a and p_b decreases by .9, the value of $P(H|E_T)$ also decreases by .9. In between these limit cases, when $0 < r < 1$, the peers should place “intermediate” weight on each other's judgments: if the value of p_a and

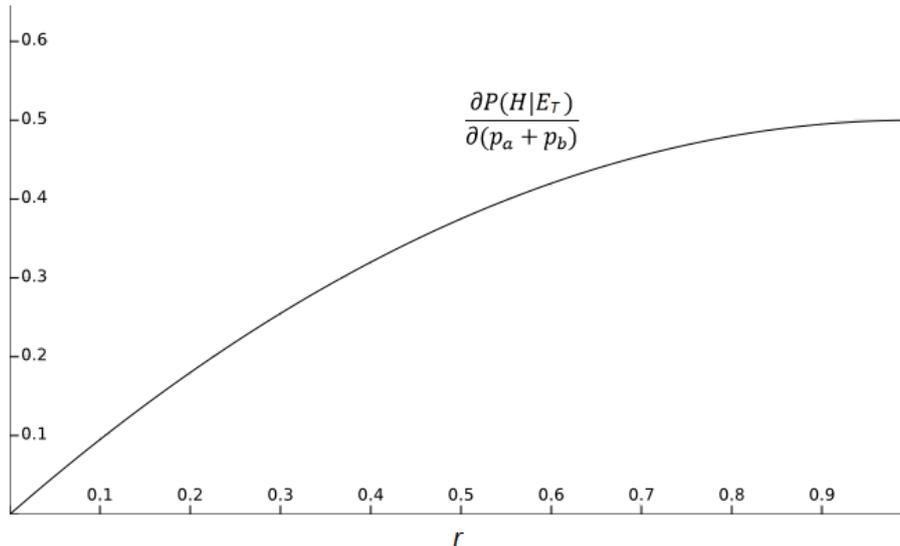


Figure 2: The derivative of $P(H|E_T)$ with respect to $p_a + p_b$ shows how r determines the impact of p_a and p_b on $P(H|E_T)$: high values of r result in a high such impact, and vice versa.

p_b changes by a given amount, the value of $P(H|E_T)$ is going to change by a somewhat smaller amount, depending on the exact value of r . For example, if $r = .6$ and we increase the average value of p_a and p_b by $.3$, the value of $P(H|E_T)$ will increase by $.3 \cdot .5(1 - (1 - .6)^2) \approx .13$.

To illustrate VEW at work, consider a case along the lines of *Weather forecasters* with which we began this section:

*Weather forecasters**: Two weather forecasters a and b assess the same body of meteorological data E and make private predictions of whether (H) it will rain later that day: a judges that $P(H|E) = .9$, and b judges that $P(H|E) = .3$. The weather forecasters are rationally certain that they both have a reliability of $r = .7$ at judging the evidence. Which credence should they adopt in H upon disagreement?

If we plug these numbers into VEW, we get the following verdict:

$$P(H|E_T) = \frac{1 - (1 - .7)^2}{2} (.9 + .3) + \frac{(1 - .7)^2}{2} \approx .59$$

So according to VEW, the weather forecasters should adopt a credence of .59 in H upon disagreement.

We now go on to explore some important consequences of VEW (§§4-5), and generalize the view to cases of non-peer disagreement (§6).

4 Against Splitting the Difference

A view that has often been associated with the equal weight view is the linear averaging or “split the difference” view:¹³

Split the Difference (SD): the parties to a peer disagreement should adopt a credence identical to the average, $.5(p_a + p_b)$, of their initial credences p_a and p_b .

An important difference between VEW and SD is that VEW takes the peers’ reliability into account, whereas SD does not. Consequently, SD yields the same verdict regardless of how reliable the disagreeing peers are, whereas VEW yields different verdicts for different values of r . For example, suppose that $r = .3$, $p_a = .8$, and $p_b = .9$. According to SD, the peers should adopt a credence of .85 upon disagreement. By contrast, VEW recommends the peers to adopt a significantly lower credence of .68 because of the peers’ low reliability. Since low values of r results in a low impact factor of p_a and p_b on $P(H|E_T)$ the peers should only adopt a credence that is slightly higher than .5 although their judgments are quite high. This is how VEW reflects the spirit behind the variably equal weight dictum.

Although VEW generally differs from SD, there are two special cases where VEW recommends splitting the difference:

$$P(H|E_T) = .5(p_a + p_b) = .5 \quad \text{for } p_a + p_b = 1.$$

¹³For discussions of linear averaging, see Jehle & Fitelson (2009), Cohen (2013), Christensen (2007a; 2009; 2011), Elga (2007), Matheson (2015), Moss (2011), Staffel (2015), and others.

$$\lim_{r \rightarrow 1} P(H|E_T) = .5(p_a + p_b)$$

The former condition represents a case where the average of the peers' judgments is .5. Here, we should expect linear averaging to be the correct response: intuitively, if two peers "pull equally hard in opposite directions," their resulting credence should not deviate from .5. The latter condition represents a limit case where the disagreeing peers are close to being perfectly reliable. Here, we should also expect linear averaging to be the correct response: when r is very high, there is a very low probability that both peers have misjudged the evidence, which means that the higher-order probability distribution $Q(p)$ will be almost fully concentrated on the peers' judgments p_a and p_b . Given this—as we show in detail in a moment—linear averaging falls out as the right response. Needless to say, two perfectly reliable peers cannot really disagree (at least given The Uniqueness Thesis), so VEW will never strictly speaking recommend splitting the difference unless $p_a + p_b = 1$. But it remains true that VEW *approaches* SD, when r *approaches* 1—that is, highly reliable disagreeing peers should approximately split the difference.

Why might one be tempted to interpret the equal weight dictum as linear averaging? We suspect that proponents of SD implicitly assume that whenever two peers a and b disagree, either (i) a has judged the evidence correctly and b hasn't, or (ii) b has judged the evidence correctly and a hasn't, but not (iii) both a and b have misjudged the evidence. Given this assumption, the equal weight dictum requires that we assign a probability of .5 to both (i) and (ii):

$$Q(p_a) = Q(p_b) = .5$$

This means that the higher-order probability distribution is fully concentrated on the peers' judgments p_a and p_b (see figure 3). If we plug this distribution into LoTP, we immediately arrive at SD:

$$P(H|E_T) = p_a Q(p_a) + p_b Q(p_b) = .5(p_a + p_b). \quad (4)$$

So SD follows straightforwardly from the assumption that whenever two peers disagree, one of them must have judged the evidence correctly. The reason why SD fails is that this assumption is false: two disagreeing peers may well both have misjudged the evidence, and the likelihood of this possibility depends on how reliable the peers are.^{14,15}

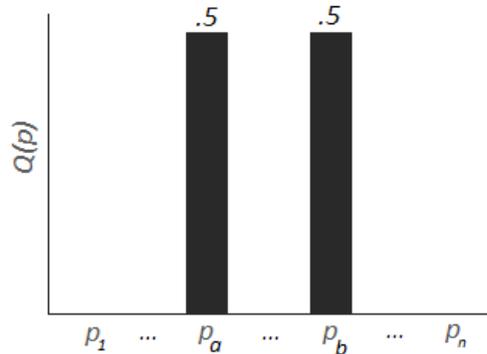


Figure 3: The higher-order probability distribution in *Peer disagreement* given the assumption that, whenever two peers disagree, one of them must have judged the evidence correctly.

While we have focused our discussion on the equal weight view in this paper, note that all views of disagreement, whether conciliatory or steadfast, should take into account the possibility that two disagreeing peers may both have misjudged the evidence. To our knowledge, SD is the only quantitatively precise view of disagreement that can be shown to fail in this regard, but the mistake is one that should be avoided by all contenders in the disagreement debate.

¹⁴An anonymous reviewer has rightly pointed out to us that the parties to a disagreement may well consider the possibility that they have both misjudged the evidence and yet end up concluding (mistakenly) that splitting the difference is the best way of revising their credences. While this is true, our claim is that SD, understood as a general *normative* view of peer disagreement, only holds under the false assumption that whenever two peers disagree, one of them will have judged the evidence correctly.

¹⁵While we have cashed out our discussion of peer disagreement in terms of credences rather than binary beliefs, it is worth noting that SD also fails in a binary belief framework. For even in a binary belief framework, two disagreeing peers may both have misjudged the evidence. Suppose, for example, that the evidence supports a belief in H , and suppose that a believes not- H , while b is agnostic about H . In this case, both a and b have misjudged the available evidence. So scenario (iii) can arise both in a credence framework and in a binary belief framework.

5 Synergy

A desirable feature of VEW is that it gives rise to plausible instances of what Easwaran et al. (forthcoming) call *synergy*—an effect whereby the parties to a disagreement should become more (or less) confident in the disputed proposition than any of them were prior to the disagreement. To illustrate, consider a case adapted from Christensen (2009, p. 759) where synergy seems warranted:

Second opinion: Two doctors independently analyze the same CT scan, and they both judge that radiation will be a more effective treatment than chemotherapy. Yet, well aware that it's easy to overlook important details on CT scans, both doctors remain open to the possibility that chemotherapy is preferable to radiation: one doctor adopts a credence of .76 that radiation is preferable to chemotherapy, and the other adopts a credence of .77. To minimize the risk of conducting the wrong treatment, the doctors decide to consult each other for a second opinion. Once they learn about each other's opinions, they both increase their confidence that radiation is preferable to chemotherapy.

Here, it seems rational for the doctors to become more confident than any of them were initially. So an adequate account of peer disagreement should, it seems, give rise to synergy in some circumstances. At the same time, it is clear that synergy is not always a rational response to disagreement. For example, if the doctors in the case above had made complete opposite judgments of the CT scan, it would be irrational of them to respond synergistically to their disagreement. Thus, an adequate account of peer disagreement should give rise to synergy in some, but not all, circumstances. In the remainder of this section, we show that VEW gives rise to synergy in a restricted set of circumstances, and argue that those are circumstances in which synergy is warranted.

As a first step, we need a precise definition of synergy. Here, we take our clue from *Second opinion* in which synergy arises in three steps. First, the doctors independently judge the CT scan. Second, each doctor adopts

what we will call her *considered credence*—that is, her credence upon having judged the CT scan and revised her credence in light of the possibility of having misjudged the evidence. Third, the doctors learn about each other’s opinions and adopt a credence that is higher than both of their considered credences prior to disagreement.

To make this characterization precise, we need a way of determining an agent’s considered credence c given that the agent has a reliability of r and has judged that $P(H|E) = j$. By reapplying, *mutatis mutandis*, the procedure used in §3 to derive VEW, we arrive at the following simple relationship between c , j , and r :

$$c = r \cdot j + \frac{1 - r}{2}. \quad (5)$$

This formula says how a single agent should take her own fallibility into account, much like VEW says how two disagreeing peers should take their fallibility jointly into account. As we should expect, c gets closer and closer to j as r approaches 1. This reflects that a reliable agent is likely to judge the evidence correctly, whereas an unreliable agents is unlikely to do so. Given this way of determining an agent’s considered credence prior to disagreement, we can go on to define synergy as follows:

Synergy: a case of peer disagreement gives rise to synergy iff $P(H|E_T) \notin [c_{min}, c_{max}]$, where c_{min} is the considered credence of the least confident peer, and c_{max} is the considered credence of the most confident peer.

If we use VEW and (5) to substitute expressions for $P(H|E_T)$, c_{min} , and c_{max} in this definition, we get the following expression:

$$\frac{1 - (1 - r)^2}{2} (j_{min} + j_{max}) + \frac{(1 - r)^2}{2} \notin \left[r \cdot j_{min} + \frac{1 - r}{2}, r \cdot j_{max} + \frac{1 - r}{2} \right].$$

This double inequality has two solutions which correspond to the conditions under which VEW gives rise to synergy (we omit the formal details):

Synergy conditions: VEW gives rise to synergy iff either

$$Syn_{<} : j_{max} < .5 \text{ and } r < \frac{2j_{max} - 1}{j_{max} + j_{min} - 1}, \text{ or}$$

$$Syn_{>} : j_{min} > .5 \text{ and } r < \frac{2j_{min} - 1}{j_{max} + j_{min} - 1}.$$

Here, $Syn_{<}$ represents the conditions under which the peers should both revise their considered credences *downwards* (i.e. $P(H|E_T) < c_{min}$). Conversely, $Syn_{>}$ represents the conditions under which the peers should boost each other's considered credences *upwards* (i.e. $P(H|E_T) > c_{max}$).

It is easily verified that the values of $(2j_{max} - 1)/(j_{min} + j_{max} - 1)$ and $(2j_{min} - 1)/(j_{min} + j_{max} - 1)$ both increase when the value of $j_{max} - j_{min}$ decreases. The value of $j_{max} - j_{min}$ intuitively represents the extent to which the peers disagree: when $j_{max} - j_{min}$ is high, the peers disagree to a high extent, and vice versa. We can thus think of VEW as giving rise to synergy under two broad sets of circumstances: either when both peers make judgments lower than .5 and their shared reliability is sufficiently low compared to the extent to which they disagree, or when both peers make judgments higher than .5 and their shared reliability is sufficiently low compared to the extent to which they disagree. Equivalently, VEW does *not* give rise to synergy when $.5 \in [j_{min}, j_{max}]$, or when the peers have a high reliability compared to the extent to which they disagree.

To motivate that *Synergy conditions* successfully tracks the circumstances under which synergy is intuitively warranted, consider the following four variations of *Weather forecasters*:

*Weather forecasters***: Two weather forecasters with a shared reliability of $r = .8$ independently analyze the same body of meteorological data, and form judgments of the evidential probability that it will rain later today.

(a) Let $j_{min} = .7$ and $j_{max} = .71$. Then $c_{min} = .66$ and $c_{max} = .67$. Here, VEW gives rise to synergy: the weather forecasters should adopt credence .7 upon disagreement, and $.7 \notin [.66, .67]$.

(b) Let $j_{min} = .29$ and $j_{max} = .3$. Then $c_{min} = .33$ and $c_{max} = .34$. Here, VEW gives rise to synergy: the weather forecasters should adopt credence .3 upon disagreement, and $.3 \notin [.33, .34]$.

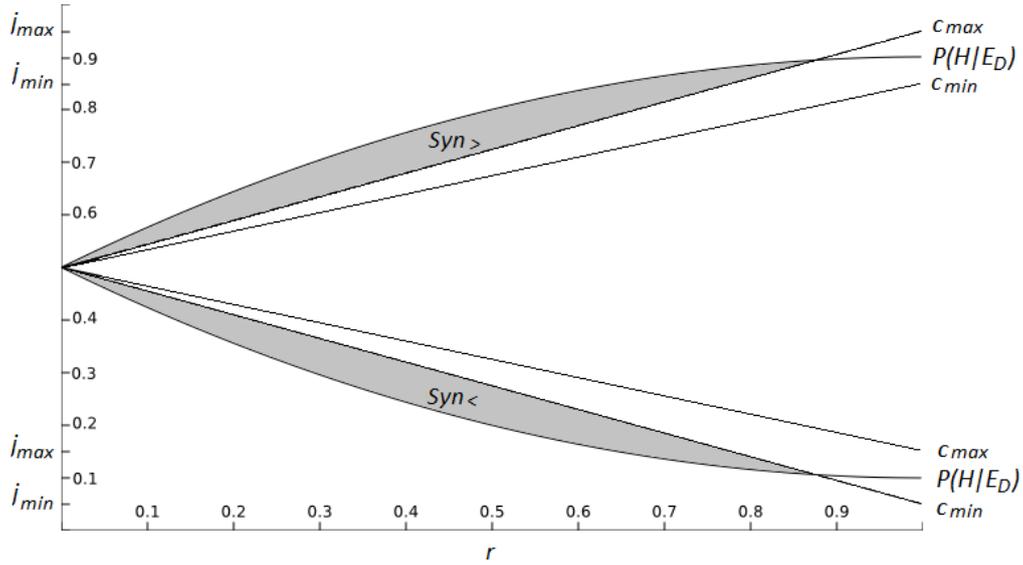


Figure 4: The shaded areas mark the conditions under which VEW gives rise to synergy. The upper part of the graph illustrates a case where $j_{min} = .85$ and $j_{max} = .95$, while the lower part of the graph illustrates a case where $j_{min} = .05$ and $j_{max} = .15$.

(c) Let $j_{min} = .3$ and $j_{max} = .8$. Then $c_{min} = .34$ and $c_{max} = .74$. Here, VEW doesn't give rise to synergy: the weather forecasters should adopt credence .55 upon disagreement, and $.55 \in [.34, .74]$.

(d) Let $j_{min} = .51$ and $j_{max} = .99$. Then $c_{min} = .59$ and $c_{max} = .86$. Here, VEW doesn't give rise to synergy: the weather forecasters should adopt credence .74 upon disagreement, and $.74 \in [.59, .86]$.

The first two cases (a) and (b) represent situations like *Second opinion* where two reliable peers only disagree to a small extent. In such cases, synergy seems warranted as two peers are less likely to be mistaken than is either one of them. The third case (c) represents situations where two peers make judgments on opposite sides of .5. In such cases, synergy does not seem warranted: if two peers pull in opposite directions, they should end up with a credence somewhere in the interval between their initial credences. Finally, case (d) represents situations where two peers have a high shared reliability compared to the extent to which they disagree. In such cases, synergy also does not seem warranted: if I have a high considered credence of .99 in H

and learn that my peer has a much lower considered credence of .51 in H , it seems that I should lower my confidence in H (and my peer should increase his). In sum, then, VEW seems to give rise to synergy under the right set of circumstances.

It is also worth noting that VEW gives rise to synergy in cases of *peer agreement*. If $j_{min} = j_{max} = j$, VEW reduces to the following view:

$$P(H|E_T) = (1 - (1 - r)^2) \cdot j + \frac{(1 - r)^2}{2}.$$

It is easy to see that $P(H|E_T) < c$ when $c < .5$, and $P(H|E_T) > c$ when $c > .5$, where $c = r \cdot j + .5(1 - r)$ is the peers' shared considered credence prior to disagreement. As such, VEW always gives rise to synergy in cases of peer agreement (except when $c = .5$): if the peers start out with a considered credence higher than .5, they should boost each other's considered credences upwards, and if they start out with a considered credence lower than .5, they should boost each other's considered credences downwards. For example, if $r = .6$ and $c = .7$, then $P(H|E_T) = .78$, and if $r = .6$ and $c = .2$, then $P(H|E_T) = .08$.

Let us close our discussion of synergy with a brief comparison to Easwaran et al. (forthcoming) who have proposed a heuristic, called *Upco*, for how to update one's credence upon learning the credence of another agent. If c_a and c_b represent the initial credences of two agents a and b , *Upco* recommends the agents to adopt the following posterior credence:

$$c_a^+ = c_b^+ = \frac{c_a c_b}{c_a c_b + (1 - c_a)(1 - c_b)} \quad (\text{Upco})$$

Upco is designed to mimic Bayesian conditionalization in cases where it is reasonable to suppose that agents have truth-tracking credences—that is, when agents are likely to have high credences in truths and low credences in falsehoods. By combining a formally precise version of this assumption with Bayes' theorem, Easwaran et al. show that *Upco* falls out as the resulting update rule.

There are two important differences between VEW and *Upco*. First, VEW is derived from an evidence-based notion of reliability, whereas *Upco* is derived from a truth-based notion of reliability. Second, VEW is supposed to apply in all cases of peer disagreement, regardless of how reliable the peers are, whereas *Upco* is only supposed to apply in cases where agents are expected to be reliable (in the truth-based sense). Despite these differences, however, we can draw a few lines of comparison between VEW and *Upco* in relation to synergy. First, *Upco* agrees with VEW that synergy should not arise when two agents initially “pull in opposite directions:” it is easily shown that $c_a^+ = c_b^+ \in [c_a, c_b]$ when $c_a \leq .5 \leq c_b$. Second, *Upco* agrees with VEW that synergy should arise when two reliable agents only disagree to a small extent. For example, if $c_a = .7$ and $c_b = .71$, *Upco* recommends that the agents adopt a credence $c_a^+ = c_b^+ \approx .85$ upon disagreement. But third, *Upco* seems to disagree with VEW when it comes to cases where two reliable agents with initial credences on the same side of .5 disagree to a high extent. For while we have seen that VEW does not give rise to synergy in such cases, it is easily verified that *Upco* gives rise to synergy *whenever* $.5 \notin [c_a, c_b]$. For example, if $c_a = .99$ and $c_b = .51$, *Upco* yields a posterior credence slightly higher than .99. This result makes sense given the assumption that agents have truth-tracking credences. For given this assumption, even credences that are only slightly higher than .5 should be taken as a slight indication that the relevant proposition is true rather than false. At the same time, however, it seems clear that two disagreeing peers with initial considered credences .99 and .51 respectively should not respond synergistically to their disagreement. Rather, when the most confident peer learns that her peer is much less confident than herself, this should be taken as an indication that her own judgment of the evidence was too optimistic to begin with. So for the purposes of capturing the conditions under which synergy should arise in paradigmatic cases peer disagreement, VEW seems to be superior to *Upco*.

6 Non-Peer Disagreement

So far we have restricted our attention to idealized cases of disagreement between epistemic peers who, by stipulation, share all of their relevant evidence and are equally reliable at judging that evidence. The benefit of such idealizations is that they make a precise analysis considerably more tractable. Yet, the interest of idealized cases of disagreement largely derives from the light they may shed on disagreement more generally. As such, it is worthwhile exploring how VEW may be generalized to cases of non-peer disagreement. Here, we will focus on three central weakenings of idealizations that are often assumed in discussions of peer disagreement:

Reliability non-peerhood: Two agents are reliability non-peers (on a given occasion) with respect to a batch of evidence E bearing on a hypothesis H iff (on that occasion) they are not equally likely to judge E correctly.

n-party disagreement: A disagreement is an n -party disagreement iff there are n parties to the disagreement.

Evidential non-peerhood: Two agents are evidential non-peers (on a given occasion) with respect to a hypothesis H iff (on that occasion) they do not possess the same total body of evidence bearing on H .

We begin by implementing the first two generalizations. Let a_1, \dots, a_n be n different evidential peers with (possibly different) reliabilities r_1, \dots, r_n . On a given occasion, the agents make (possibly different) judgments p_1, \dots, p_n of the evidential probability of H given E . Given this, we can reapply, *mutatis mutandis*, the procedure from §3 to arrive at the following generalization of VEW (we omit the formal details):

$$P(H|E_T) = \left(1 - \prod_{i=1}^n (1 - r_i)\right) \frac{\sum_{i=1}^n p_i r_i}{\sum_{i=1}^n r_i} + .5 \left(\prod_{i=1}^n (1 - r_i)\right). \quad (6)$$

As we should expect, (6) reduces to VEW when $n = 2$ and $r_1 = r_2$. Moreover, it is easily verified that (6) recommends the agents to adopt a credence identical to the average of their judgments under two special circumstances:

when the average of their judgments is .5, or when $r_1 = \dots = r_n \rightarrow 1$. Under these conditions, (6) in effect reduces to a generalized version of SD.

Like VEW, (6) is based on our idealized notion of reliability according to which an agent has a reliability of r just in case there is a probability of r that the agent makes a perfect judgment of the evidence. But as previously mentioned, agents like you and I presumably rarely (if ever) judge evidence so precisely. It seems more reasonable to think of an agent's reliability as the likelihood that the agent's judgment will at least be "in the right ballpark." A natural way to make this idea precise is to model an agent's reliability in terms of a more or less spread-out higher-order probability distribution. Suppose that agent a judges that $P(H|E) = p_a$ on a given occasion. We can then let $\sigma_a(p)$ be a higher-order probability function representing what we will call a 's *reliability distribution* on the given occasion. On this model, a has a high reliability just in case the value of $\sigma_a(p)$ is high for values of p close to p_a , and low for values of p far from p_a . Given that the probability variable p is continuous, σ_a is a density function which should be normalized as usual ($\int_0^1 \sigma_a(p) dp = 1$).

Given this setup, we can then use LoTP to calculate the first-order probability of H given the conjunction of the first-order evidence E and the higher-order evidence E_a given by a 's judgment:

$$P(H|E_a) = \int_0^1 p \cdot \sigma_a(p) dp. \quad (7)$$

This equation yields an agent's considered credence given our new model of reliability, just as (5) yields an agent's considered credence given the old model of reliability.

Suppose, next, that n different evidential peers a_1, \dots, a_n have (possibly different) reliability distributions $\sigma_1, \dots, \sigma_n$ on a given occasion. To determine which credence the agents should adopt upon disagreement, we need to combine $\sigma_1, \dots, \sigma_n$ to arrive at a final higher-order probability distribution which we can call σ^+ . Given σ^+ , we can use LoTP to determine $P(H|E_T)$, where, as usual, E_T is the total body of evidence consisting of the conjunction of E

and all of the agents' judgments:

$$P(H|E_T) = \int_0^1 p \cdot \sigma^+(p) dp. \quad (8)$$

The combined higher-order probability distribution σ^+ is obtained by normalizing the product of $\sigma_1, \dots, \sigma_n$:

$$\sigma^+(p) = \frac{\prod_{i=1}^n \sigma_i(p)}{\int_0^1 \prod_{i=1}^n \sigma_i(p) dp}.$$

By substitution in (8), we arrive at the following generalization of VEW:

$$P(H|E_T) = \int_0^1 p \cdot \frac{\prod_{i=1}^n \sigma_i(p)}{\int_0^1 \prod_{i=1}^n \sigma_i(p) dp} dp. \quad (9)$$

To illustrate this view, suppose that $n = 2$, and suppose that the two agents' reliability distributions are given by the following normalized beta distributions (see figure 5):

$$\begin{aligned} \sigma_a(p) &= 30 \cdot p \cdot (1-p)^4 \\ \sigma_b(p) &= 105 \cdot p^4 \cdot (1-p)^2. \end{aligned}$$

If we plug these reliability distributions into (7), we get a and b 's considered credences prior to disagreement:

$$\begin{aligned} P(H|E_a) &= \int_0^1 p \cdot 30 \cdot p \cdot (1-p)^4 dp \approx .29 \\ P(H|E_b) &= \int_0^1 p \cdot 105 \cdot p^4 \cdot (1-p)^2 dp \approx .63. \end{aligned}$$

If, instead, we plug the reliability distributions into (9), we get the credence that a and b should adopt upon disagreement:

$$P(H|E_T) = \int_0^1 p \cdot \frac{30 \cdot p \cdot (1-p)^4 \cdot 105 \cdot p^4 \cdot (1-p)^2}{\int_0^1 30 \cdot p \cdot (1-p)^4 \cdot 105 \cdot p^4 \cdot (1-p)^2 dp} dp \approx .46$$

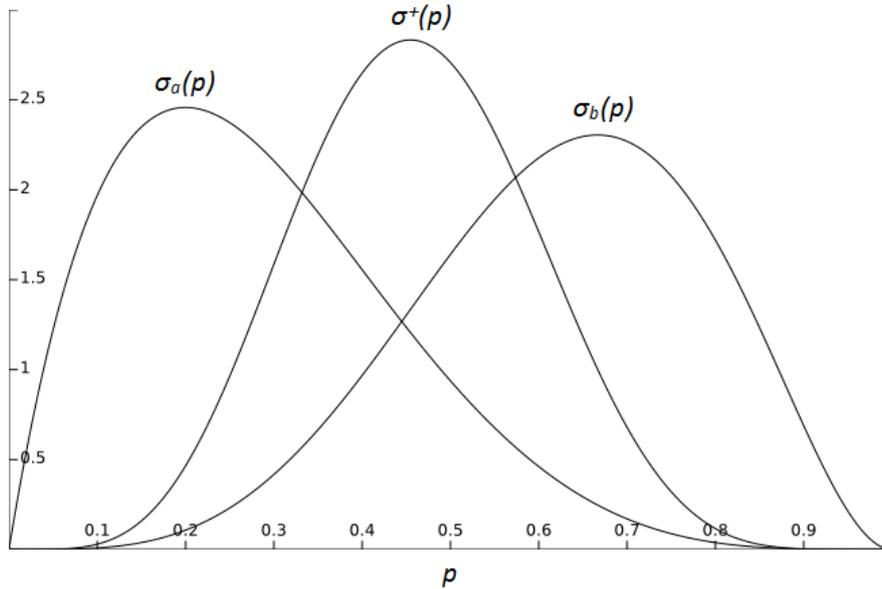


Figure 5: The reliability distributions of two agents a and b on a given occasion of disagreement.

Note that a and b should not respond synergistically to their disagreement ($.46 \in [.29, .63]$) which was to be expected since a and b initially pull in opposite directions.

While VEW generalizes quite straightforwardly to cases of disagreement between n agents with (possibly) different reliabilities, it is less clear how to generalize VEW to disagreements between evidential non-peers. As Christensen (2007a, pp. 211-212) notes, we may reasonably suppose that evidential *parity* produces the same results as evidential *peerhood*: it doesn't seem to make a difference whether the disagreeing parties have the same evidence, or whether they have different, but equally good, evidence. But what happens when the disagreeing parties do not possess equally good evidence? To answer this question, we need a quantitative measure of the relative quality of different bodies of evidence. This measure should influence the weight that should be placed on the opinions of different agents: the better an agent's evidence, the more weight should (other things being equal) be placed on the agent's opinion. We are optimistic that our framework can eventually be generalized to capture this idea. But we leave a detailed investigation for

future work.

7 Conclusion

We have outlined a simple formal framework for reasoning about the evidential impact of a body of higher-order evidence, and have used this framework to motivate a novel interpretation of the equal weight view of disagreement—we have called it the Variably Equal Weight View (VEW). According to VEW, disagreeing peers should always place equal weight on each other’s opinions, but the amount of weight they should give each other depends on how reliable they are: roughly, reliable peers should give each other equal *high* weight, while unreliable peers should give each other equal *low* weight. We showed that VEW differs considerably from the standard Split the Difference (SD) interpretation of the equal weight view, and showed that VEW gives rise to plausible instances of synergy. Lastly, we explored how VEW may be generalized to cases of non-peer disagreement.

Acknowledgments. We would like to thank David Christensen, Mikkel Gerken, Thomas Grundmann, Bjørn Hallsson, Klemens Kappel, Saber Malmgren, Jon Matheson, Nikolaj Nottelmann, Josefine Pallavicini, and an anonymous referee for *Episteme* for helpful comments and discussion. Earlier drafts of this paper were presented at the 2015 workshop “The Epistemology of Disagreement” at Copenhagen University, the 2016 Annual Meeting of the Danish Philosophical Society, and the 2016 NordPhil Conference at Uppsala University; we thank the audiences at these occasions for valuable feedback.

References

- Ballantyne, N. and E.J. Coffman (2011). “Uniqueness, Evidence, and Rationality”. In: *Philosophers’ Imprint* 11.18.
- (2012). “Conciliationism and Uniqueness”. In: *Australasian Journal of Philosophy* 90.4, pp. 657–670.
- Christensen, D. (2007a). “Epistemology of Disagreement: The Good News”. In: *Philosophical Review* 116.2, pp. 187–217.

- (2009). “Disagreement as Evidence: The Epistemology of Controversy”. In: *Philosophy Compass* 4.5, pp. 756–767.
- Christensen, D. (2010). “Higher-Order Evidence”. In: *Philosophy and Phenomenological Research* 81.1, pp. 185–215.
- (2011). “Disagreement, Question-Begging, and Epistemic Self-Criticism”. In: *Philosophers’ Imprint* 11.6.
- (2014). “Conciliation, Uniqueness and Rational Toxicity”. In: *Noûs* 49.4.
- Cohen, S. (2013). “A Defense of the (Almost) Equal Weight View”. In: *The Epistemology of Disagreement: New Essays*. Ed. by D. Christensen and J. Lackey. OUP, pp. 98–117.
- Easwaran, K. et al. (forthcoming). “Updating on the Credences of Others: Disagreement, Agreement, and Synergy”. In: *Philosophers’ Imprint*.
- Elga, A. (2007). “Reflection and Disagreement”. In: *Noûs* 41.3, pp. 478–502.
- (2010). “Subjective Probabilities Should Be Sharp”. In: *Philosophers’ Imprint* 10.05.
- Feldman, R. (2006). “Epistemological Puzzles About Disagreement”. In: *Epistemology Futures*. Ed. by Stephen Hetherington. Oxford University Press, pp. 216–236.
- (2007). “Reasonable Religious Disagreements”. In: *Philosophers Without Gods: Meditations on Atheism and the Secular*. Ed. by Louise Antony. OUP, pp. 194–214.
- (2009). “Evidentialism, Higher-Order Evidence, and Disagreement”. In: *Episteme* 6.3, pp. 294–312.
- Fitelson, B. and D. Jehle (2009). “What is the “Equal Weight View’?” In: *Episteme* 6.3, pp. 280–293.
- Greco, D. and B. Hedden (forthcoming). “Uniqueness and Metaepistemology”. In: *Journal of Philosophy*.
- Inwagen, P. Van (2010). “We’re Right. They’re Wrong”. In: *Disagreement*. Ed. by Richard Feldman and Ted A. Warfield. OUP Oxford.
- Jaynes, E. (1973). “The Well-Posed Problem”. In: *Foundations of Physics* 3.4, pp. 477–493.
- (2003). *Probability Theory: The Logic of Science*. Cambridge University Press: Cambridge.

- Kelly, T. (2005). “The Epistemic Significance of Disagreement”. In: *Oxford Studies in Epistemology, Volume 1*. Ed. by J. Hawthorne and T. Gendler. Vol. 1. OUP, pp. 167–196.
- (2010). “Peer Disagreement and Higher Order Evidence”. In: *Social Epistemology: Essential Readings*. Ed. by A. Goldman and D. Whitcomb. Oxford University Press, pp. 183–217.
- Keynes, J. (1921). *A Treatise on Probability*. Dover Publications.
- Lackey, J. (2008a). “What Should We Do When We Disagree?” In: *Oxford Studies in Epistemology*. Ed. by Tamar Gendler and John Hawthorne. Vol. 3. OUP, pp. 274–93.
- Lackey, J. (2008b). “A Justificationist View of Disagreement’s Epistemic Significance”. In: *Proceedings of the Xxii World Congress of Philosophy*. Ed. by Alan Millar Adrian Haddock and Duncan Pritchard. Vol. 53. OUP, pp. 145–154.
- Lasonen-Aarnio, M. (2014). “Higher-Order Evidence and the Limits of Defeat”. In: *Philosophy and Phenomenological Research* 88.2, pp. 314–345.
- (2015). ““I’m Onto Something!” Learning About the World by Learning What I Think About It”. In: *Analytic Philosophy* 56.4, pp. 267–297.
- Levinstein, B. (2013). “Accuracy as Epistemic utility”. Rutgers University.
- (2015). “With All Due Respect: The Macroepistemology of Disagreement”. In: *Philosophers’ Imprint* 15.13, pp. 1–20.
- Matheson, J. (2009). “Conciliatory Views of Disagreement and Higher-Order Evidence”. In: *Episteme: A Journal of Social Epistemology* 6.3, pp. 269–279.
- (2015). *The Epistemology of Disagreement*. Palgrave.
- Moss, S. (2011). “Scoring rules and epistemic compromise”. In: *Mind* 120.480, pp. 1053–1069.
- Pettigrew, R. (2016). “Accuracy, Risk, and the Principle of Indifference”. In: *Philosophy and Phenomenological Research* 92.1, pp. 35–59.
- Schoenfield, M. (2013). “Permission to Believe: Why Permissivism Is True and What It Tells Us About Irrelevant Influences on Belief”. In: *Nouûs*, pp. 193–218.

- (2015). “A Dilemma for Calibrationism”. In: *Philosophy and Phenomenological Research* 91.2, pp. 425–455.
- Staffel, J. (2015). “Disagreement and Epistemic Utility-Based Compromise”. In: *Journal of Philosophical Logic* 44, 273–286.
- Titelbaum, M. (2015). “Rationality’s Fixed Point (Or: In Defense of Right Reason)”. In: *Oxford Studies in Epistemology* 5, pp. 253–294.
- White, R. (2005). “Epistemic Permissiveness”. In: *Philosophical Perspectives* 19.1, pp. 445–459.
- Whiting, D. (2016). “Against SecondOrder Reasons”. In: *Noûs* 49.4, pp. 1–23.
- Worsnip, A. (2014). “Disagreement About Disagreement? What Disagreement About Disagreement?” In: *Philosophers’ Imprint* 14.18.