

A Pragmatic Argument against Equal Weighting

Ittay Nissan-Rozen & Levi Spectre

Abstract

We present a minimal pragmatic restriction on the interpretation of the weights in the “Equal Weight View” (and, more generally, in the “Linear Pooling” view) regarding peer disagreement and show that the view cannot respect it. Based on this result we argue against the view. The restriction is the following one: if an agent, i , assigns an equal or higher weight to another agent, j , (i.e. if i takes j to be as epistemically competent as him or epistemically superior to him), he must be willing – in exchange for a positive and certain payment – to accept an offer to let a completely rational and sympathetic j choose for him whether to accept a bet with positive expected utility. If i assigns a lower weight to j than to himself, he must not be willing to pay any positive price for letting j choose for him. Respecting the constraint entails, we show, that the impact of disagreement on one’s degree of belief is not independent of what the disagreement is discovered to be (i.e. not independent of j ’s degree of belief).

0. Introduction

How do rational agents respond to disagreement with their peers? This seemingly simple question triggers controversy even among those who frame it in Bayesian terms, a framework that has proven instructive in the study of rationality. The correct response to peer disagreement is controversial, among other reasons, because it is entangled with several difficult questions and issues: What is peerhood? What kind of evidence justifies identifying an agent as one’s peer? What is the proper relation between higher and lower order rational credence (roughly, how should having equal confidence in the adequacy of one’s own response to evidence and another agent’s response influence one’s first order belief)? Is it okay to count shared data as evidence against another agent’s competence in responding to this data?

This paper attempts to make some headway on the question of disagreement without engaging directly with these formidable issues. Our main contribution takes the form of a pragmatic constraint on the notion of peerhood: if an agent, j , is your peer, then assuming that j is sympathetic – she wants you to gain as much as possible – you should be willing – in exchange for a certain payoff – to let her decide for you whether to accept a bet with positive expected utility. If you are not willing to accept this exchange even for a sure

payoff, you do not seriously regard *j* as your peer. We also generalize this constraint to cover all cases of *partial* epistemic deference, not only the case of peerhood.

The constraint clearly conflicts with views –“Steadfast” views¹ – that allow you to disregard a peer’s degree of belief in a disputed proposition, *P*, or to regard her as less likely than yourself to have the correct degree of belief in *P* after discovering the disagreement.

Yet the pressure here on Steadfast views is not significant since proponents of such views already admit that they only regard a peer as being equally likely to get things right before her degree of belief is revealed, *but not after*. So, although the constraint seems to shine an unfavorable light on Steadfast views, it does so, so to speak, on a bullet already bitten.

Views that are expected to embrace our constraint are “Conciliatory” views that distinctly reject the idea that you can discount a peer’s view by comparing it to your own degree of belief (relative to the available evidence).² In particular, views that advise convergence (or near convergence) would accept our constraint after disagreement is disclosed. This is because when convergence is reached there is no relevant difference that would justify rejecting a sure gain. A prominent Conciliatory view, the “split the difference” view, requires peers to converge on the average of the difference in degree of belief that they find between them. This is the most natural and straightforward interpretation of (arguably) the most prominent Conciliatory view, which is known as “the equal weight” (or EW, for short) view.³ The constraint we offer seems to fit the EW view – particularly on its split the difference interpretation – nicely, not only in its convergence result but also in its rejection of the idea that one can change one’s view about whether one’s peer is equally likely as oneself to get things right in a particular case by consulting one’s own reasoning from the evidence.

Surprisingly, the EW view on the split the difference interpretation fails to meet the constraint we offer. Moreover, our argument applies equally well to EW views that at least sometimes – but not always – require that peers split the difference and converge (or come close to converging).⁴

¹ See, e.g. Christensen 2011 for a characterization of this view.

² Christensen 2011, p. 1-2 locates this as the central issue dividing Conciliatory and Steadfast views.

³ We thank an anonymous referee for another journal for helping us see the nuances here.

⁴ The “splitting the difference” (averaging over the difference in degrees of belief) interpretation of the EW view is defended by Elga (2007), at least as understood by Christensen (2007, p. 199, n. 15) and others (see footnote 10), and by Christensen (2011) (with some qualifications) and more recently by Cohen (2013): “when peers discover they disagree, each should adopt the simple average of their credences, i.e. they should split the difference between their credences. If I am at .8 and my peer is at .2, then we should each move to .5.”

Since our argument applies to any splitting the difference *case*, it also applies to the Conciliatory view developed by Christensen in his more recent (2009) and (2011) publications. Though Christensen distances his view from the splitting the difference interpretation of the EW view (2009, pp. 758-9), he advises peers to converge in some cases (and within a limited range of degrees of belief). The argument can apply even after taking into consideration the

More generally, our argument shows that any instance of the "Linear Pooling" view (LP view) – i.e. any view according to which after learning the degree of belief of *j*, *i*'s degree of belief would be a *weighted* average of his original degree of belief and *j*'s degree of belief – cannot respect a generalization of the constraint stated above. The generalization is the following one. If *i* assigns a higher weight to *j*, (i.e. if *i* takes *j* to be epistemically superior to him), he must be willing – in exchange for a positive and certain payment – to accept an offer to let a completely rational and sympathetic *j* choose for him whether to accept a bet with positive expected utility. If *i* assigns a lower weight to *j*, he must not be willing to pay any positive price for letting *j* choose for him.

We argue that any plausible interpretation of the weights in the LP view (of which the EW view is a special case) must respect this constraint. Thus, since the LP view violates this constraint, it must be rejected. The conclusion we wish to draw is that, in contrast with a commitment to the LP view, the effect that an advisor's belief should have on one's own belief, partially depends on what that belief ends up being. In other words, the rational impact of an advisor at any level of authority is not uniform. This makes room for the development of a new Conciliatory view that calls for varying weights of impact that depend on advisors' degrees of belief.⁵ Perhaps such a view has already been proposed, by e.g., Kelly (2010); the Total Evidence view.⁶

The remainder of this paper is organized in the following way. Section 1 discusses the relation between the EW view and LP, presents the suggested constraint in a precise way, states the main theorem that stands at the heart of our argument, and demonstrates the intuition behind it (the full proof is in the appendix). Section 2 explains why advocates of LP views must accept the constraint and discusses the relation between our argument and other arguments against the LP view one can find in the literature. Finally, section 3 discusses possible objections to our argument.

rather daunting list of obstacles in articulating a principle for when averaging of this sort is required (2009, 766, note 11). Moreover, as we will make clear, our argument does not rely on exact averaging, in fact, it merely relies on "Linear Pooling," i.e. on a view that attaches a fixed weight (e.g. 1/3) to a peer's degree of belief within a range. So our argument also seems to apply to Christensen's (2011) (e.g. p. 3, note 3) view that takes an even bleaker view regarding the exact impact a peer's view should have at the present stage of the discussion on rational disagreement (2011, p. 17). His remarks there suggest that he expects the impact of a peer's degree of belief to be less than splitting the difference but still significant, i.e. suggesting that it is fixed. Our argument, however, goes against even these vague general expectations. This is because respecting our constraint will entail, at the very least, different levels of impact to different degrees of belief that a peer may be discovered as having.

⁵ We are not committed, however, to the idea that the resulting view will be Conciliatory in Christensen's sense. Such a view may very well violate his independence constraint. See note 2 above.

⁶ The Total Evidence view seems to be committed to a rejection of the constraint we offer. Kelly (2010) endorses the split the difference view in cases where peers lose access to their original evidence (when the evidence is purely *psychological* in his sense—see note 20 below). There seems to be, however, no straightforward way for the Total Evidence view to escape our argument in such a case. But perhaps there is an externalist type of interpretation of the view that can do this by marginalizing the role of access to the evidence. This possibility will not be explored here. Thanks here to an anonymous referee.

1. The EW view, the LP view and the Exchanging Epistemic Identity Principle

Two rational agents i and j have the same body of evidence E . Their attitudes regarding propositions are represented by their respective credence (or degree of belief) functions; $c_i(\bullet)$ and $c_j(\bullet)$.⁷ The credence function of at least one of these agents, let us suppose it is agent i , is defined over propositions of the form “ j has degree of belief x that P ,” which we label X (or X -type propositions). We will use “ x ” to refer to the value of j ’s degree of belief that P (i.e. “ X ” refers to “ $c_j(P)=x$ ”). Before discovering the disagreement, then, there are expectations associated with all X propositions in a partition consisting of j ’s possible credences for P (according to i).

We assume that after i learns one of the X propositions to which $c_i(\bullet)$ assigns a positive credence value, i ’s new degree of belief in P , $c_i'(P)$, is equal to i ’s prior conditional degree of belief in P given X . We assume, that is, that $c_i'(P) = c_i(P|X)$. By using this assumption, we are following the policy, here and below, of characterizing LP views (and the EW view in particular) within the strictures of the standard Bayesian framework. We discuss possible objections to this assumption in section 3. However, it is important to stress that (like many others) we are not claiming that this is how rational agents reason, but merely that reasoning that is *necessarily* inconsistent with the Bayesian framework pays a theoretical price that needs to be counterbalanced by a special argument.⁸

The main question that arises in the setting just presented is whether it is possible to identify any compelling restrictions on $c_i(P|X)$, (i.e. on i ’s conditional degree of belief in P given that j believes P to degree x). Intuitively—and we take this to be a major merit of the LP view—the value of $c_i(P|X)$ depends on i ’s assessment of j ’s epistemic competence. The more competent (with regards to propositions such as P) i takes j to be (in a given context), the closer $c_i(P|X)$ will be to x . i ’s assessment of j ’s competence can be represented as i ’s commitment to a function that specifies for every possible ordered pair of x and $c_i(P)$ (i.e. for every combination of i ’s initial degree of belief in P and j ’s degree of belief in P), what impact j ’s degree of belief will have on $c_i(P|X)$. “Linear Pooling” is the name given in the literature to one possible (but prominent) restriction on this function:

⁷ Throughout the paper we will assume that E is shared between i and j . This will allow us to investigate the role epistemic competence plays in cases of epistemic disagreement in isolation from informational considerations.

⁸ See Steele (2012) and Lasonen-Aarnio (2013) for two statements of the kind of attitude (with some minor differences) that we prescribe to. The demand that a rational agent’s response for discovering a disagreement will be consistent with Bayesian updating is also discussed in Jehle and Fitelson (2009) and Shogenji (2007). We return to this issue in section 3.

Linear Pooling (LP): $c_i(P|X) = wx + (1-w)c_i(P)$, for some $0 < w < 1$.

In words, i 's conditional degree of belief in P given j 's degree of belief in P is a weighted average of the two agents' degrees of belief. The more competent i takes j to be, the higher the weight he assigns to j 's degree of belief (we discuss the interpretation of the weights more thoroughly below).

The important point about this formula is that the weights chosen are independent of j 's degree of belief that P . i needs to commit himself to a formula that specifies his conditional degree of belief in P given that j believes P to degree x , *which holds for any x* . This is exactly the claim against which we argue in this paper.⁹ What is the alternative? Well, the alternative we argue for in this paper is to simply reject this restriction on the way rational agents learn from others. We do not argue here for any alternative restriction. However, it is obvious that a natural guideline for the development of an alternative is the following one: the weight agent i gives to agent j 's opinion should decrease as the distance between x (j 's degree of belief in P) and $c_i(P)$ (i 's initial degree of belief in P) increases. We are not going to defend a worked-out suggestion, though. Possibly there is no rational impact algorithm to defend. In this paper we only show that any view that incorporates (restricted) linear pooling – of which the EW views are instances – is flawed.

Now, one special case of LP – a case that has attracted considerable attention in the literature – concerns the case in which i takes j to be his *peer*, i.e. i takes j to be as epistemically competent (with regards to P -type propositions) as he takes himself to be. In such a case, it is tempting to demand that the weight i assigns to himself will be equal to the weight he assigns to j . This is naturally, though not necessarily, understood as the EW view on its split the difference interpretation.¹⁰ Assigning equal weight to your peer need not mean that you should split the difference with her, but it has been taken to mean that, at least within a limited range

⁹ At least within a certain range of values x might take. It seems that Christensen (2007), a prominent split the difference view advocate at the time (but see footnote 4 above), agrees that when distance is substantial enough, applying the equal weight principle is not necessarily warranted. Such is the case, for instance, when a peer is presenting an obviously “crazy” view about relatively simple matters. In more recent writing, Christensen specifies another case where if your own credence is high and so is your peer's, you should perhaps be more confident that the proposition in question is true even if your peer's degree of belief is slightly lower than yours (Christensen 2009, 759). Yet for other cases, like the one he specifies just before giving this counterexample, within a range at least, he advocates splitting the difference (Ibid. 758-9).

Since our argument works against views that restrict the LP formula to a range of possible x values, Christensen's counterexamples – as far as we can see – do not motivate a view that goes far enough away from LP. We see no reason, however, why a view such as Christensen's could not incorporate the idea that, even if the discovery of a disagreement should not influence the degree to which one takes an advisor to be rational, the impact of this view should not be constant (even within a range of values).

¹⁰ In addition to advocates of the EW view mentioned in footnote 3 and 6, Jehle and Fitelson (2009) and Kelly (2010) formulate the EW as the split the difference view, as do also Enoch (2010, p. 2) and Wilson (2010) as an interpretation of Elga's (2007) EW view. In contrast, Lasonen Aarnio (2013) discusses the relation between the EW view – that one assigns the same probability to being correct as a peer being correct – and splitting the difference. She contends that the relation between the two depends on transparency of the agent's second order attitudes.

of values, the impact of the disagreement is at least uniform and usually (at least close to) the average of the difference between the peers (i.e. half in the LP schema above). Here we will follow this usage and focus *as an example* on the case where the weights do entail averaging (or splitting the difference). Nothing in our argument will depend on this choice, as we show below.

The Equal Weight (EW) view: When i takes j to be as epistemically competent as himself (regarding the matter at hand): $c_i(P | "c_j(P) = x") = \frac{c_i(P) + x}{2}$

Although the argument in this paper applies to the LP view generally, since some of the discussions in the literature target specifically the case in which i takes j to be his epistemic peer and interpret this as averaging, at some points in our discussion it will be convenient to concentrate on the EW view as stated above. This is done solely for the purpose of clarity of presentation and nothing in our argument is restricted to this specific case. On the contrary, as will become clear there is a sense in which the argument is even stronger when it comes to weights different from 0.5.

We wrote earlier, informally, that the weights in the LP view express something like i's assessment of j's level of epistemic competence with regard to P. When i takes j to be his peer, he takes j to be as competent as himself with regards to P-type propositions. What, however, does it mean to take another agent to be "as epistemically competent as oneself" with regard to a certain class of propositions? In the literature no clear answer to this question can be found beyond vague demands such as that if i takes j to be his peer he must take j to be as likely to "get things right" as i takes himself to be.

This, of course, does not say much until we know what "getting things right" means in the context that the EW view advocate is envisioning.¹¹ One strength of the argument we offer here is that it does not rely on any specific interpretation of the weights in the LP view. Rather, it applies to any interpretation of the weights that respects a very minimal pragmatic constraint which, we believe, any EW advocate should endorse. Here it is:

¹¹ While the answer to this question is straightforward in a non-Bayesian framework ("getting things right" with respect to a proposition, P, is to believe P in case P is true and believe $\neg P$ in cases it is false), in a Bayesian framework it is not. Lasonen-Aarnio (2013) suggests that "getting things right" should be understood in terms of identity between one's subjective probability and the "correct" evidential (or "epistemic") probability. She then shows that such an understanding leads to some problematic results for the EW view. This is, however, far from being the only possible way to interpret the phrase. Another interesting direction is to understand "peerhood" in terms of getting equal "scores" in some measure of accuracy (see for example Anders Levinstein 2015), and there are also other possible directions. One of the advantages of the approach we take here is that we do not have to commit ourselves to any single interpretation of peerhood. Instead we use our minimal pragmatic constraint that any advocate of the EW view must accept (regardless of the way she understands the term "peerhood").

*Exchanging Epistemic Identity (EEI)*¹²

Suppose *i* faces the following choice: *i* can either **accept** a bet, or **decline** it or **pass** the choice of whether to accept it or not to *j*. Suppose *i* is certain that

a) *j* is rational (in the decision theoretic sense)

and

b) *j* is completely sympathetic to *i* in the following sense: if *i* lets *j* choose for him, *j* maximizes expected utility when the utility used is *i*'s utility function, but the probability function used is *j*'s credence function.

Then:

1. If *i* assigns to *j* a weight, w , such that $w \geq 0.5$, *i* should be willing to pass the choice of whether to accept the bet to *j* **in exchange for a certain positive payment**.
2. If *i* assigns to *j* a weight, w , such that $w < 0.5$, *i* should not be willing to **pay any certain positive payment** in order to pass the choice of whether to accept the bet on to *j*.

In the next section we discuss the EEI more extensively. However, the intuition behind it is clear. Since by stipulation *i* is certain that *j* is rational and uses *i*'s utility function, the only difference between *i* making the choice and *j* making it is the probability function used.

Now, it does not make sense for *i* to pay a price just so that the choice will be taken using a probability function that *i* himself views as an inferior one, and it does not make sense for him to decline a certain payoff just so that the choice will be taken using a probability function that he himself views as a superior one. In particular, EW view advocates (in peer disagreement situations) are expected to consider this a kind of deciding-myself-fetishism.

Thus, we believe any plausible interpretation of the weights in the LP view must respect the EEI. However consider the following theorem:

¹² We thank an anonymous referee for suggesting the name.

Theorem 1

For any credence function of i that assigns a non-trivial probability value to the possibility that j 's degree of belief in P is different from i 's degree of belief in P , and for any non-trivial weight, w :¹³

- a. Even if $w > 1/2$ (i.e. even if i takes himself to be more epistemically competent than j), there always exists a bet such that i will be willing to pay a positive amount of utility in exchange for letting j choose for him whether to accept this bet.

- b. Even if $w \leq 1/2$ (i.e. even if i takes himself to be less epistemically competent than j), there always exists a bet such that i will be willing to pay a positive amount of utility in exchange for avoiding passing the choice of whether to accept the bet to j .

The theorem shows that independently of i 's credence distribution over the partition of hypothesis about j 's degree of belief in P (with one – uninteresting – qualification: the case in which i is *certain* that j holds exactly the same degree of belief as he does), and for every non-trivial weight, it is *always* possible to construct a bet such that i will violate the highly intuitive restriction on the interpretation of the weights we have introduced.

It is interesting that when w is trivial the result no longer holds. This makes sense: to assign a weight of 0 to an agent (possibly to oneself) is to ignore that agent's opinion – and if one ignores an agent's opinion there can be no dependency between that opinion and the weight assigned to the agent.

Our argument is, then, the following one. The theorem shows that the LP view violates the EEI. However, the EEI must be respected by any plausible interpretation of the weights in the LP view. Hence, there is no plausible interpretation of the weights in the LP view and thus the LP view is false. In the next two sections we better defend that EEI and discuss several possible objections, but before doing so it may be helpful to gain a better understanding of the theorem.

¹³ We prove the theorem only for the case in which i assigns a positive probability to a countable number of hypothesis about j 's degree of belief in P . This is so since it is unclear how the LP view should be defined in the case of uncountably many such hypothesis (due to the fact that in such cases each hypothesis typically gets 0 probability). In any case, if the LP view is correct it surely must be correct also for the countable case and so proving the theorem only for this case is enough. We discuss this a little bit further in the appendix.

The proof of the theorem is in the appendix. However, the intuition behind the proof can be easily demonstrated using a simple example that involves the case of equal weights. Consider the following bet. If P, i gets 1 unit of utility; if $\neg P$ i gets 0 units of utility. i has to decide between:

- a. Accepting the bet.
- b. Rejecting the bet and getting a certain 0.2 units of utility.
- c. Passing the choice whether to accept the bet to j.

Suppose further that i takes j to be his peer, that i believes P to degree 0.5 and that i assigns a positive and equal credence value of 0.5 to only two hypothesis about j's degree of belief in P:

X – “j believes P to degree 0.9”

Y – “j believes P to degree 0.1”

Notice that so far all is good: the values specified together with the peerhood assumption and the LP formula are in accord with the Theorem of Total Probability:

$$1. \quad c(P) = 0.5 = c(P|X)c(X) + c(P|Y)c(Y) = (w_i c(P) + 0.9w_j)c(X) + (w_i c(P) + 0.1w_j)c(Y) = (0.5 \times 0.5 + 0.5 \times 0.9)0.5 + (0.5 \times 0.5 + 0.1 \times 0.5)0.5 = 0.5$$

Notice now that i's expected utility for a (accepting the bet) is 0.5 and his expected utility for b (rejecting the bet) is 0.2. What is i's expected utility for c (passing the choice to j)?

Well, if Y is true (if j's degree of belief in P is 0.1) j will choose to reject the bet. In this case, i gets 0.2 units of utility. If X is true (if j's degree of belief in P is 0.9) j will choose to accept the bet. In this case i will face a bet with expected utility of

$$2. \quad c(P|X) \times 1 + c(\neg P|X) \times 0 = c(P|X) = c(P) \times 0.5 + x \times 0.5 = 0.5 \times 0.5 + 0.9 \times 0.5 = 0.7.$$

Thus, i's expected utility from c is

$$3. \quad EU(c) = 0.5 \times 0.2 + 0.5 \times 0.7 = 0.45 < 0.5$$

In other words, i's expected utility from c is lower than i's expected utility from a and so i will violate the EEI: he will be willing to pay a positive amount in order to avoid passing the choice of whether to accept the bet to j even though i takes j to be as epistemically competent as himself.

In this example, like in the general proof, we take advantage of the fact that if *i* would pass the choice of whether to accept the bet to *j*, *j*'s choice would reveal to *i* her degrees of belief in *P* and thus when calculating the expected utility of passing the choice to *j*, *i* must take this under consideration (*i* must, that is, use his conditional degrees of belief given *j*'s possible choices).

Thus, on the pragmatic level, i.e. when it comes to calculating expected utility, the way *i* treats *j*'s opinion does depend on what this opinion turns out to be (because *j*'s beliefs are related to her choices in a systematic way) and this pragmatic dependency is translated using the EEI to an epistemic dependency.

2. The EEI and other arguments against the LP view

In the literature one can find several other arguments against the LP view (and several arguments against the EW view on its averaging interpretation). A popular – and quite successful – *Bayesian* line of argument against the LP view is based on the observation that it runs into trouble when moving from the simple case that involves a single updating after learning the degree of belief of a single agent in a single proposition to the more general case that involves successive updating on different types of evidence that include several other people's degrees of belief in several different propositions together with some first order evidence.

For example, Shogenji (2007) shows that there are cases in which a Bayesian agent cannot update (using classical Bayesian updating) his beliefs after learning both some first order propositions and the degrees of belief of another agent in some propositions in a way that obeys the LP constraint without violating the laws of probability.¹⁴ Bradley (2007) shows that even when it comes to a single updating instance, the LP view leads to problematic consequences when more than one epistemic peer is involved. For example, he shows that in some cases a Bayesian agent can respect the LP constraint with regard to a proposition, *P*, only if he treats the degree of belief in *P* of one of the agents to which he assigns a positive weight as independent of the *truth of P*.

Although we are sympathetic to these arguments, we do not take them to refute the LP view. First, there are some potential replies. For example, Jehle and Fitelson (2009) show that there are several variations of the EW view (they only deal with the split the difference EW view, not with the LP view in general) that are invulnerable to the problem discussed by Shogenji (2007) and Wilson (2010). Similarly, Steel (2012) shows that there is an interesting class of cases in which the problem to which Bradley (2007) points does not arise.

¹⁴ Wilson (2010) discusses the same mathematical phenomenon from a different philosophical angle.

Second, and more generally, the existing arguments against the LP view: (1) Do not hold in the case of a single updating on the degree of belief of a single agent in a single proposition; (2) Only show that *there are cases* in which the LP view runs into trouble. They do not show that the LP view always runs into trouble. Thus, these arguments allow for an acceptance of the LP view *in a restricted class of cases*.

As a matter of fact, there is an extensive body of literature (that started growing in the beginning of the 80s) that discusses the problem of probabilistic opinion pooling using an axiomatic approach.¹⁵ This literature identifies several different sets of axioms that uniquely pick out different methods of pooling. The LP method, for example, is the only pooling method that respects the following two axioms: (1) Unanimity Preservation: if all the agents to which *i* assigns a positive weight have the same degree of belief in *P* as *i*, *i* should not change her degree of belief in *P* after learning the other agents' degrees of beliefs; (2) Independence: *i*'s degree of belief in *P*, after learning the other agents' degrees of belief, depends only on the other agents' degrees of belief in *P* (not on their degrees of belief in other propositions).¹⁶

Now, since these sets of axioms are inconsistent with each other, and since each one of the sets seems to be the appropriate one to adopt in different contexts, it seems clear that no single pooling method should be pointed to as the right method to adopt *in every context*. Rather, it seems that different methods should be used in different contexts.¹⁷ Arguments such as those of Shogenji (2007), Wilson (2010) and Bradley (2007) should be understood, we believe, as pointing to specific contexts in which we should not expect Unanimity preservation and Independence to hold. However, these arguments do not show that the LP approach should *never* be used.

Our argument, on the other hand, does exactly that. Our argument: (1) Works in the case of a single updating on the degree of belief of a single agent in a single proposition; (2) Can be made against *any such instance* of LP that involves non-trivial weights (i.e. weights that are different from either 1 or 0).

Our argument can do this because, unlike the arguments mentioned above, it does not point to any problematic *epistemic* features of the LP view. Rather it rests on the very weak *pragmatic* restriction on the interpretation of the weights in the LP view we have introduced. The power of our argument depends, then, entirely on the plausibility of the EEI *from the point of view of advocates of the LP view*.

As explained in section 1 the main idea behind the EEI is that, on the assumption that *j* is rational and sympathetic to *i* (which *we stipulate* is believed to degree 1 by *i*), the only difference between *j* choosing

¹⁵ See List and Dietrich 2014 for a good introduction.

¹⁶ This was proven first by Lehrer and Wagner (1981).

¹⁷ Thus, after discussing several sets of axioms, List and Dietrich argue that "it should be clear that there is no one-size-fits-all approach to probabilistic opinion pooling" (List and Dietrich 2014, p. 20).

whether to accept the bet and i choosing is which credence function is used in the expected utility calculation, i's or j's. If the weights express i's estimation of the relative level of epistemic competence of himself vs. j (with respect to P), then violation of the EEI amounts to being willing to pay in order to make a choice using a credence function that i himself judges to be inferior.

From this perspective, the EEI can be seen as a natural extension of similar restrictions regarding the conceptual inter-dependence between desires, beliefs and actions. Arguably, it is *conceptually impossible* for a **rational** agent to prefer “x units of utility if A and 0 if $\neg A$ ” to “x units of utility if $\neg A$ and 0 if A” while believing A to a lower degree than $\neg A$. In the same way – if “epistemic competence” is a meaningful term – it seems to be *conceptually impossible* for a **rational** agent to (strictly) prefer using one probability function, $c(\bullet)$ to using another probability function $c'(\bullet)$ in expected utility calculations while taking $c'(\bullet)$ to be more epistemically competent than $c(\bullet)$.

When it comes to the case of peerhood, this consideration is conceptually related to the idea of symmetry. One thing the EW view on the averaging interpretation (in particular) has going for it, is that it is symmetric: if i and j are peers (and know this) they should treat each other in the exact same way. This symmetry is – at least according to advocates of the EW view – dictated by the concept of peerhood. When two rational agents are peers, identity does not matter: if a given reaction to the disagreement is rational for one of them, it is also rational for the other. In the case of the choice presented to the agent in our EEI condition, if it is rational for i to accept (or to decline) the offer to let j choose for him whether to accept a bet (knowing that j adopts i's utility function when making the decision), it is also rational for j to accept (decline) the offer to let i choose for her whether to accept the bet (knowing that i adopts j's utility function when making the decision). Thus, in the case in which i and j have the same utility function, there is no difference whatsoever between j choosing and i choosing and so a violation of the EEI in the case of equal weights amounts to a refusal to get a free payoff (which is clearly conceptually impossible when the payoff is in terms of utility). This is so, it is important to stress again, in the EW view's advocate own light.¹⁸

Another reason the EW view (specifically) justifies EEI is that assuming it provides the rational response to disagreement and that i knows j is rational, after credences are exposed, i will surely respect the EEI. This is because on the EW view, if i and j are rational, they will converge — they will have the same credence for P and hence, whatever payoff EEI offers will be gained for free. In rejecting the EEI the EW view advocate would have to explain why before discovering the disagreement j's credences are not to be trusted as much as i's,

¹⁸ Thanks to an anonymous referee for the suggestion to justify the EEI by symmetry considerations. This referee also suggested that there are interesting connections between the EEI and some discussions in the literature, e.g., how Keren (2007) understands deference in the context of testimonial knowledge transfer. Unfortunately, discussing possible connections would take us too far afield.

while after the discovery they are good enough to be an equal basis for the convergence. It is hard to see how this feat can be accomplished.¹⁹ The argument we have employed here is novel in this sense. It focuses on the epistemic constraints imposed by conditionalization and the LP view as employed *before* disagreement is discovered. This feature of our argument blocks any attempt to evade it by pointing to a particular property of the disagreement itself.

In contrast to advocates of the LP view, rejecting the EEI does seem to be a plausible option *for those who reject linear pooling*: if the weight one gives to one's peer is a function of that peer's credence value, then it makes sense that one will be willing to pass the choice of whether to accept a bet with a positive expected utility to one's peer only conditional on that peer holding certain credence values. In other words, rejection of linear pooling opens the door to a view according to which it is not possible to decide with complete certainty whether someone is a peer (in the sense of giving her credence a certain weight) regarding P before learning what this person's degree of belief in P is.

Proponents of linear pooling cannot, however, make this move because they believe that once one has decided to treat another person as one's peer with respect to P, that person's degree of belief in P will have the same impact no matter what it turns out to be (at least within a range of values). This is the essence of the LP view: its commitment to the independence of the weight assigned to an agent from that agent's degree of belief.

Some scholars understand, however, the requirement of the independence of the weights in a way that is slightly different from the one presented here. Instead of taking the debate to be about whether the weight that i gives j's opinion depends on j's *degree of belief* that P, they take it to be about whether the weight that i gives to j's opinion depends on the *reasoning* j employs when adopting her degree of belief in P (see for example Christiansen 2009 and Elga 2007).²⁰

We do not suggest that these two formulations of the debate are equivalent. They are not. However, in cases (like those we consider here) in which i and j share the same evidence, an important part of what we call "j's

¹⁹ One way, though, was hinted at by a referee of this journal. Suppose we employ Kelly's (2010) distinction between the discovery of one's peer's reaction to the evidence—psychological evidence—and non-psychological evidence—the original evidence. The EW advocate can claim that without the psychological evidence, a peer's credence is not to be relied on.

We don't find this line particularly plausible, but we don't want to claim that no development in this direction will work. Kelly himself uses the distinction to argue *against* the EW view. Roughly, his claim is that it is implausible that psychological evidence swamps non-psychological evidence. See note 6 above for the claim that Kelly's Total Evidence view is susceptible to our argument for some cases, cases for which the possible reply on behalf of the EW view will be equally problematic.

²⁰ Here is Christensen's (2011) formulation: "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."

reasoning" is the impact the evidence has on j 's degree of belief that P . Although there might be some mental activities that should be qualified as "reasoning" and cannot be captured nicely in our setting as either instances of identifying the evidence or as instances of deciding what the impact of the evidence on one's degree of belief should be, we ignore this option. By doing so we only set the bar higher for ourselves: we show that even when the only difference between j 's and i 's reasoning is what they take to be the impact of the evidence (that they share) on their degree of belief that P , the weight i assigns to j should *not* be independent of j 's reasoning.

Thus, we conclude any advocate of the LP view should be committed to the EEI and thus subject to our argument.

3. Possible objections

We have shown that the LP view is inconsistent with the EEI. In the previous section we have argued that any advocate of the LP view must be committed, however, to the EEI. Thus, the conclusion must be that the LP view is false. How else can a LP view advocate respond?

One possible response is to reject our setting by arguing that it, in some way, misrepresents the LP view. We are not going to respond here to the general charge that the Bayesian framework cannot capture *examples* of peer disagreement according to the LP view. We take that to be a significant theoretical cost in and of itself. The LP view is not (and we think it should not be) viewed as opposed to Bayesianism. But there are more specific claims that can be used to resist conditionalization within the Bayesian framework in peer disagreement situations.

A seemingly attractive way to resist the argument above is to claim that in our example, at least, i is certain in advance that j has a different credence than he has (though he knows not which and cannot change his credence). Thus he knows that j disagrees with him. The EW view (on any of its interpretations) concerns the reaction to the *discovery* that a peer disagrees, so perhaps we are misrepresenting the debate.

Attractive as this response may seem, it won't do. It concerns our simplified example. It makes no difference to the argument, if we add a third X-type proposition, say Z (where Z ="j's degree of belief is 0.5"), such that $c_i(Z)=0.8$ (and adjusting i 's degree of belief in X and Y accordingly symmetrically since X , Y , and Z form a partition). The theorem shows that the resulting case(s) where certainty of disagreement is absent (and in fact i is almost certain that there is no disagreement) will violate the EEI in the same way the example does.

In general it is interesting to note, though, that in contrast with the cases presented in the literature, cases such as the one in our example show that the discovery of disagreement itself is not an essential element of the problem of peer disagreement. What does seem essential is the discovery of what the disagreement is, what j 's degree of belief actually is, not merely that j 's degree of belief is different from i 's. If there is a rational response to the discovery of peer disagreement, it seems possible to construct cases where the possible values are balanced in such a way that an agent will know that she disagrees with her peer but will not thereby be required to change her credence before discovering what the disagreement is.

Related to the previous objection is the claim that the credence function $c(\bullet)$ should not contain X -type propositions at all (i.e. hypothesis about j 's degree of belief in P). The motivation for this is that the assumption that agents always start from a prior probability distribution that is defined over an algebra that contains all X s is so demanding that it makes any conclusion that is based on it practically insignificant. A Bayesian alternative is to argue that, upon learning agent j 's degree of belief in P , i updates his degree of belief in P in a way that gives the required LP view value, and then after this initial change he updates his degrees of belief in other propositions using Jeffrey conditionalization. The problem to be investigated according to this line of thought is how the initial updating should go. Thus the question of peer disagreement is not as we have posed it, but rather: *how should agent i update his degree of belief in P after learning X having no prior probability distribution that is defined over an algebra that contains X to refer to?*

While it seems to us that this alternative line of representing disagreements is more focused on disagreement and more realistic, we also believe that using our previous presentation of the LP view (i.e. in terms of conditional credences) carries no significant theoretical cost. What motivates moving from our way to the alternative is the understanding that actual agents have limited cognitive resources and thus cannot always anticipate their belief updates for learning each one of the propositions that it is possible for them to learn (and specifically propositions such as X). This is one motivation for considering pooling strategies in general (as opposed to discussing what the conditional credences should be). Thus, when real people learn propositions such as X they usually do not have a prior conditional probability for P given X to refer to. We should still demand, however, that any good answer to the question regarding real agents and their limitations *be consistent* with the right answer to the question of disagreement as we have framed it. Agents could, after all, *sometimes* at least, have the required conditional degrees of belief.

We should demand, that is, that the rational way i updates his beliefs after learning X in case $c(\bullet)$ is defined over an algebra that does not contain X will be consistent with the way i would update his degree of belief in P in a case where $c(\bullet)$ is defined over an algebra that does.

Our response to this attempt to avoid the argument is also how we would want to answer worries in the same general vicinity. In general, we can get to the LP view values as a response to peer disagreement either by a non-conditionalization rule or by an anti-conditionalization rule. We have no problem with non-conditionalization rules as such, but we think they do not avoid the argument. After adopting such a rule one would still need to show how the argument is avoided with regard to cases that overlap with conditionalization. We do have a problem, however, with anti-conditionalization rules. In fact, as much as we like our EEI constraint, we think adopting an anti-conditionalization rule is a higher theoretical price to pay than rejecting it simply because it doesn't accord with the LP view. That is, simply biting the bullet is better here.

To conclude, what seems clear is that linear pooling – the independence of the weight from the agent's degree of belief – has unacceptable consequences precisely because of this independence. Unless you assign a weight of 1 or 0 to the opinion of another agent, learning what this other agent's opinion is must effect – contra the EW view – the weight you assign to her opinion. What exactly should this effect be? This, unfortunately, is still an open question.

Appendix

Theorem 1

For any credence function of i that assigns a non-trivial probability value to the possibility that j 's degree of belief in P is different from i 's degree of belief in P , and for any non-trivial weight, w :

- 1. Even if $w > 1/2$ (i.e. even if i takes himself to be more epistemically competent than j), there always exists a bet such that i will be willing to pay a positive amount of utility in exchange for letting (a completely rational and sympathetic) j choose for him whether to accept this bet.*
- 2. Even if $w \leq 1/2$ (i.e. even if i takes himself to be less epistemically competent than j), there always exists a bet such that i will be willing to pay a positive amount of utility in exchange for avoiding passing the choice of whether to accept the bet to j .*

Proof:

Assume $0 < c(P) < 1$ and let us rescale i 's utility function so that the payment i receives by choosing to accept the bet in case P is true is 1, and the payment i receives by choosing to accept the bet in case P is false is 0.

Thus, i 's expected utility from accepting the bet is $c(P)$. If i chooses to reject the bet he gets a payment of $0 < a < c(P)$. Thus, i prefers accepting the bet to rejecting it.

Now let $Y = \text{"j's degree of belief in } P \text{ is lower than or equal to } a\text{"}$. Since j is rational and sympathetic, in case Y is true (i.e. in case j 's degree of belief in P is lower than or equal to a), j will choose to reject the bet. Similarly, in case $\neg Y$ is true (i.e. in case j 's degree of belief in P is higher than a), j will choose to accept the bet.

Let L be the act of letting j choose whether to accept the bet, and let $\neg L$ be the act of rejecting the offer to let j choose whether to accept the bet. We saw that i 's expected utility from $\neg L$ is just the expected utility of the bet, which is $c(P)$. We can now rewrite $c(P)$ using the Theorem of Total Probability applied to the partition $\{Y, \neg Y\}$:

$$4. \quad EU(\neg L) = c(P) = c(P|\neg Y)c(\neg Y) + c(P|Y)c(Y)$$

What is the expected utility of L ?

$$5. \quad EU(L) = c(P|\neg Y)c(\neg Y) + ac(Y)$$

In words: i 's expected utility from passing the choice on to j equals the probability i assigns to the possibility that j will accept the bet multiplied by the expected utility of the bet from the point of view of i after learning that j chooses to accept the bet (i.e. $c(P|Y)$) plus the probability i assigns to the possibility of j rejecting the bet multiplied by the expected utility of rejecting the bet (i.e. a).

i will choose to pass the choice on to j iff $EU(\neg L) \leq EU(L)$. Thus, it immediately follows from 4 and 5 that if $c(Y) > 0$, i chooses to pass the choice on to j iff $a > c(P|Y)$.

Let Y^* be the set of all propositions of the form " j 's degree of belief in P is x_i " such that $x_i \leq a$, to which i assigns a positive probability. In other words, Y is the disjunction of all the propositions in Y^* . We assume that Y^* is countable. We suspect that a similar theorem to the one we present here can be proved for the uncountable case (this becomes clear by considering the rest of the proof). However, our characterization of the LP view is not well-defined in the uncountable infinite case (in which typically each proposition gets a 0 probability) and we believe that the theoretical cost involved in the philosophical interpretation of a characterization which is mathematically well-defined also in the uncountable case outweighs the benefits associated with such a characterization. This is so since a position according to which LP is true *only* in the case in which i assigns positive probabilities to uncountably many hypothesis regarding j 's degree of belief in P , seems unmotivated. If LP is true, it must be true also in the countable case (indeed, also for the finite

case), it seems to us. Thus, proving the theorem for the countable case is enough for the philosophical goal we have set for ourselves.

Now, since, by definition Y is the disjunction of all the propositions in Y^* (we assume, of course, that all the propositions in Y^* are disjoint, i.e. that j 's degree of belief in P is unique):

$$6. \quad c(P|Y) = \frac{\sum_{X_i \in Y^*} c(P|X_i)c(X_i)}{c(Y)} = \frac{\sum_{X_i \in Y^*} c(P|X_i)c(X_i)}{c(Y)}$$

However, from LP, for each X_i , $c(P|X_i) = wc(P) + (1-w)x_i$. Thus:

$$7. \quad c(P|Y) = \frac{\sum_{X_i \in Y^*} c(X_i)(wc(p) + (1-w)x_i)}{c(Y)} = wc(p) + (1-w) \frac{\sum_{X_i \in Y^*} c(X_i)x_i}{c(Y)}$$

Let us now define $Z = \frac{\sum_{X_i \in Y^*} c(X_i)x_i}{c(Y)}$ and we get:

$$8. \quad c(P|Y) = wc(p) + (1-w)Z$$

Notice that Z is a weighted average of all x_i in Y^* and thus is lower than a (or equal to it in the case in which Y^* includes only one proposition: " j believes P to degree a "). Similarly $c(P|Y)$ must be strictly lower than $c(P)$.

We are interested in the expression " $a - c(P|Y)$ ". We want to show that for some values of a such that $c(P) > a > 0$, and $c(Y) > 0$, this expression is positive and for some values it is negative *independently of w* . If this is the case then, independently of the weight, there are always bets such that $EU(-L) > EU(L)$ and bets such that $EU(L) > EU(-L)$.

$$\underline{a - c(P|Y) < 0}$$

Case 1: if there is an x^* which is the lowest value for j 's degree of belief in P to which i assigns a positive probability, set $a=x^*$ (so that $c(Y) > 0$). From the LP formula we get:

$$9. \quad c(P|Y) = wc(P) + (1-w)a$$

and since $a < c(P)$, it immediately follows that $c(P|Y) > a$ and so $a - c(P|Y)$ is negative.

Case 2: If there is no x^* which is the lowest value for j 's degree of belief in P to which i assigns a positive probability, then pick an a such that $a < w_i c(P)$, and then (from equation 8):

$$10. c(P|Y) = wc(P) + (1 - w)Z > a$$

and trivially $a - c(P|Y) < 0$ and $c(Y) > 0$ (since there is no lowest x^*).²¹

$a - c(P|Y) > 0$:

Let V be the propositions "j's degree of belief in P is lower than $c(P)$ " and let V^* be the set of all propositions of the form "j's degree of belief in P is x_i " such that $x_i < c(P)$ (i.e. V is the disjunction of all the propositions in V^*). Let us now define $Z^* = \frac{\sum_{X_i \in V^*} c(X_i)x_i}{c(V)}$. Notice that Z^* is a weighted average of all x_i in V^* and thus is strictly lower than $c(P)$.

Now set:

$$a = wc(P) + (1 - w)Z^*$$

With such an a , it follows from equation 8 that

$$11. a - c(P|Y) = (1 - w)(Z^* - Z)$$

However, since $a < c(P)$, $(Z^* - Z) \geq 0$.

Case 1: there is a X_i in V^* such that $x_i > wc(P) + (1 - w)Z^*$. In such a case clearly $(Z^* - Z) > 0$ and thus $a - c(P|Y) > 0$.

Case 2: there is no such X_i . In such a case there exists a X_i in V^* with the highest x_i of all the x_i s in all the X_i s in V^* (in other words there exists a hypothesis about j 's degree of belief in P to which i assigns a positive probability that assigns to j the maximal degree of belief in P which is strictly lower than $c(P)$). Let us call this x_i , x^{**} . By the definition of Y , $c(P|Y)$ is equal for all values of a as long as $a > x^{**}$ and so it is always possible to find a high-enough a such that $c(P) > a > c(P|Y)$. In this case $c(Y) > 0$ is trivially true.

²¹ We thank an anonymous referee for suggesting this move.

This concludes the proof.

References

Anders-Levinstein, Benjamin (2015), With All Due Respect: The Macro-Epistemology of Disagreement, *Philosophers' Imprint*, 15 (13), 1-20.

Bradley, R. (2007), Reaching a Consensus, *Social Choice and Welfare*, 29, 609–632.

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

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

Cohen, Stewart (2013), A Defense of the (Almost) Equal Weight View, In David Christensen & Jennifer Lackey (eds.), *The Epistemology of Disagreement*.

Dietrich, F. and List, C. (2014), Probabilistic Opinion Pooling, in *Oxford Handbook of Probability and Philosophy*, Forthcoming.

Elga, A. (2007), Reflection and Disagreement, *Nous*, XLI(3):478–502.

Jehle, D., and Fitelson, B. (2009). What is the "Equal Weight View"? *Episteme*, 6:3. 280–93.

Kelly, Thomas (2010). Peer disagreement and higher order evidence. In Alvin I. Goldman & Dennis Whitcomb (eds.), *Social Epistemology: Essential Readings*. Oxford University Press. pp. 183--217.

Keren, A. (2007). Epistemic Authority, Testimony and the Transmission of Knowledge. *Episteme*, 4(3), 368-381.

Lackey, Jennifer & Christensen, David (eds.) (2013). *The Epistemology of Disagreement: New Essays*. Oxford University Press.

Lasonen-Aarino, M. (2013), Disagreement and Evidential Attenuation, *Nous*, 47:4,767-94.

Lehrer, K and Wagner, C. (1981), *Rational consensus in science and society*, Dordrecht: Reidel.

Lewis, D. (1980): A Subjectivist's Guide to Objective Chance, in R. C. Jeffrey (ed.), 1980, *Studies in Inductive Logic and Probabilities*, Vol. II, Berkeley: University of California Press, pp. 263-93.

Shogenji, T. (2007), My way or her way: A Conundrum in Bayesian Epistemology of Disagreement. Unpublished manuscript.

Steele, K. S. (2012), Testimony as evidence: more problems for linear pooling, *Journal of Philosophical Logic*, 41:6, 983-99.

Wilson, A. (2010), Disagreement, equal weight and commutativity, *Philosophical Studies*, 149:3, 321-6.

Acknowledgments

Earlier versions of this paper were presented in the Haifa workshop on rational belief and normative commitment (2017), *Stockholm Higher Seminar*, Department of Philosophy, Stockholm University, Sweden November (2015). We thank the participants of these events for useful discussions. We especially want to thank David Enoch, Zeev Goldschmidt, Noam Nisan, Orri Schneebaum, and two anonymous referees for their very useful comments and suggestions as well as an anonymous referee from another journal. Levi Spectre's research was supported by the Israeli Science Foundation (Grant No. 463/12).