

An Optimality-Argument for Equal Weighting^[*]

Christian J. Feldbacher-Escamilla

Spring 2020

Abstract

[1543] There are several proposals to resolve the problem of epistemic peer disagreement which concentrate on the question of how to incorporate evidence of such a disagreement. The main positions in this field are the *equal weight view*, the *steadfast view*, and the *total evidence view*. In this paper we present a new argument in favour of the equal weight view. As we will show, this view results from a general approach of forming epistemic attitudes in an optimal way. By this, the argument for equal weighting can be massively strengthened from reasoning via epistemic indifference to reasoning from optimality.

Keywords: epistemic peer disagreement, equal weight view, steadfast view, higher order evidence, epistemic optimality, meta-induction

1 Introduction

Two peers have an epistemic disagreement regarding a proposition if their epistemic attitudes towards the proposition differ. So, e.g., an agent's degrees of belief in p might be different from that of her peer, etc. The question of how to deal with such a disagreement is the problem of epistemic peer disagreement.

Several proposals to resolve this problem have been put forward in the literature. Most of them concentrate on the question of if, and if so, to what extent one should weight evidence of such a disagreement in forming an epistemic attitude towards a proposition. A classical position is the so-called *conciliatory view* which calls for incorporating such evidence (cf., e.g., Christensen 2007; Elga 2007; Feldman 2007). A position at the other end of the spectrum is the so-called *steadfast view* which suggests to generally give no weight to such evidence (cf., e.g., Rosen 2001). In between are views that suggest to sometimes

^[*][This text is published under the following bibliographical data: Feldbacher-Escamilla, Christian J. (2020). "An Optimality-Argument for Equal Weighting". In: *Synthese* 197.4, pp. 1543–1563. DOI: [10.1007/s11229-018-02028-1](https://doi.org/10.1007/s11229-018-02028-1). All page numbers of the published text are in square brackets. For more information about the underlying project, please have a look at <http://cjf.escamilla.academia.name>.]

give such evidence some weight, and sometimes not, [1544] or to weight such evidence from case to case differently (cf., e.g., *the total evidence view* in Kelly 2011).

In this paper we want to present a new argument to the debate which has lasted for more than a decade now (it was initiated by Feldman 2007; but the systematic discussion of epistemic disagreement started much earlier; Lehrer and Wagner 1981, e.g., is central). In favour of the most prominent conciliatory view, namely the *equal weight view*, we show that equally weighting one's peer's and one's own epistemic attitudes allows for an *optimal* resolution of epistemic disagreement, whereas the other views lack this feature.

Our approach to epistemic peer disagreement is based on the theory of meta-induction as introduced by Schurz (2008). The theory of meta-induction was designed to provide an optimality argument for induction. However, Schurz (2012) indicates also an application of the theory of meta-induction to the debate on *fundamental* disagreement, where epistemic agents "disagree in their underlying *cognitive system* [... i.e. they disagree on] fundamental principles of reasoning that determine the criteria for justification". There it is suggested to resolve such disagreements by applying methods that are "universal in the sense of being reasonable in every cognitive system" (Schurz 2012, p.343 and p.346). As we show, this suggestion can be also expanded to the general case of epistemic peer disagreement and is even decisive regarding the single positions in the debate. The main idea of our application is roughly as follows:

- (i) In case of epistemic peer disagreement agents have the same inferential skills and share all evidence.
- (ii) The equal weight view suggests that in such a case one should weight all epistemic attitudes equally, whereas its competitors suggest to deviate from such a weighting.
- (iii) We suggest interpreting or operationalising *inferential skills* by help of *past accuracy* (reliability): Two agents, who share all evidence, have the same inferential skills, if they were equally accurate in the past.
- (iv) Now, according to the theory of meta-induction, if one considers a sequence of competing epistemic attitudes, then there is a procedure of weighting the attitudes in such a way that one is guaranteed to become optimal: By weighting an epistemic attitude accordingly with its past accuracy in the sequence (i.e. with its reliability), such a weighted epistemic attitude cannot be outperformed in terms of accuracy in the long run, so it is optimal in the long run.
- (v) Hence, according to this method, if the past accuracies (reliabilities) are equal, then they also should be weighted equally.
- (vi) Hence, in case of epistemic peer disagreement, the equal weight view instantiates an optimal meta-inductive weighting method, whereas its competitors do not and can become suboptimal.

- (vii) This result is robust in the sense that it holds also for less idealised model assumptions.

The plan of our investigation is as follows: In section 2 we present the formal framework of the classical approaches to epistemic peer disagreement (regarding i–ii). Subsequently, in section 3, we expand the framework to the meta-inductive setting and suggest interpreting the condition about equal inferential skills in terms of reliabilities [1545] (regarding iii). In section 4 we provide our optimality argument in favour of the equal weight view. There we also show how the other approaches fail to deal with the optimality argument (regarding iv–vi). In section 5 we list some possible objections to our argument and explain how one can overcome them (regarding vii). We conclude in section 6.

2 Approaches to Epistemic Peer Disagreement

Epistemic peer disagreement with respect to a proposition p is that special case of disagreement, where two epistemic peers differ in their epistemic attitudes towards p . Regarding the notion of *epistemic attitudes*, we follow the line of argumentation of Kelly (2011) and assume without further ado attitudes on the cardinal scale, i.e. degrees of belief. Also regarding the notion of *epistemic peers* we presuppose the well-known interpretation and assume in accordance with Feldman (2007) (and Elga 2007, Christensen 2007) that peers are characterised by their having the same evidence and the same inferential skills. Regarding evidence, we also follow Christensen (2010) and Kelly (2011, p.194) in distinguishing between *higher order evidence* which is any kind of information about the degrees of belief of an epistemic agent, and *first order evidence* which is all other information not about degrees of belief of an epistemic agent.

Now, the traditional approaches to the problem of epistemic peer disagreement differ along the line of how to incorporate higher order evidence. If we spell out ‘incorporate’, as is often suggested, in terms of linear weighting (cf. Elga 2007; Nissan-Rozen and Spectre 2017), then we can describe the case of peer disagreement as one of finding a correct weighted update of one’s degrees of belief given some higher order evidence. Note that there are other ways of weighted updating as, e.g., geometrical weighting, however, we will focus on linear characterisations of the three classical approaches to epistemic peer disagreement only (for details on geometrical weighting cf. Dietrich and List 2016; Easwaran et al. 2016; Brössel and Eder 2014).

So, our model of epistemic peer disagreement basically states that higher order evidence about such disagreement is taken into account by linear weighting. Since we want to keep track of the accuracy of the different agents at different occasions, we time-index everything – the credences, the higher order evidence, the weights (and later on also the propositions in question). For modelling this case we introduce the following notation:

- Let us assume that $Pr_i^1, Pr_i^2, Pr_i^3, \dots$ ($1 \leq i \leq n$) are the agents’ degrees of belief after updating on first order evidence – the superscript indices

mark the rounds of update.

- In case of epistemic peer disagreement between individuals i and j regarding some p it holds at some round t : $Pr_i^t(p) = r_i \neq r_j = Pr_j^t(p)$. Furthermore, the agents become aware of this – they receive in the same round but in a second phase higher order evidence e^t which contains the details of the disagreement. Thus, in this setting the problem of epistemic peer disagreement can be formulated as follows: [1546]

$$Pr_i^t(p|e^t) = Pr_i^t(p|Pr_1^t(p) = r_1, \dots, Pr_n^t(p) = r_n) = ?$$

- As we have stated above, incorporation of such higher order evidence is often described as a form of linear weighting. So it holds:

Epistemic Peer Disagreement:

There are w_1^t, \dots, w_n^t such that:

$$1. Pr_i^t(p|e^t) = \sum_{j=1}^n w_j^t \cdot Pr_j^t(p) \quad (\text{EPD})$$

This is our model of the problem of epistemic peer disagreement: A group of $1, \dots, n$ epistemic peers updates on higher order evidence e^t by determining weights w_1^t, \dots, w_n^t .

In this model, the weights can vary among the rounds. We will see later on that this is crucial for motivating our optimality-argument. In most investigations the weights are considered to be constant, i.e. round-independent. However, this expresses the model assumption that the factor of inferential competence does not change, something we want to be flexible on. Note also that the assumption of shared evidence is “hard-coded” in the model: All agents update at the same round on the same first order evidence. This is due to the fact that in our discussion we need no flexibility regarding differences in evidence.

Given this framework, we can describe three classical approaches to the problem of epistemic peer disagreement with the help of the following specifications: The equal weight view claims that in case of peer disagreement the epistemic attitudes of all peers should get equal weight:

Equal Weight View:

Among peers the weights are equal: 1. of (EPD) and:

$$2. w_1^t = \dots = w_n^t = 1/n \quad (\text{for all } 0 < t \in \mathbb{N}) \quad (\text{EWV})$$

Considering the impact of first order evidence and higher order evidence in an agent’s forming of an epistemic attitude, it is easy to see that higher order evidence can swamp first order evidence, simply because an agent’s first order evidence gains only $1/n$ weight, whereas higher order evidence gains $n - 1/n$ weight. The most prominent proponents are Christensen (2007) and

Elga (2007); a coarse-grained version for a nominal scale is held by Feldman (2007).

The steadfast view points in the completely opposite direction: According to it, one's own epistemic attitude gains full weight, whereas the attitudes of one's peers get no weight, i.e. higher order evidence does not matter at all:

Remain Steadfast View:

Among peers one's own position gets full weight: 1. of (EPD) and

$$2. w_i^t = 1 \text{ and } w_j^t = 0 \quad (\text{RSV})$$

(for all $0 < t \in \mathbb{N}, j \in \{1, \dots, n\} \setminus \{i\}$)

[1547] Perhaps the most prominent proponent of (RSV) is Rosen (2001).

Finally, we also want to model the total evidence view: According to this view, only taking into account either higher order evidence or first order evidence provides no adequate response to the total evidence available. Although the description of this view in (Kelly 2011, p.201) does not automatically ask for some *linear* "interaction" between first and higher order evidence, the arguments and further phrasing of it (cf., e.g., *swamping, insubstantial and substantial evidence, equally strong pieces of evidence, greater proportion of our total evidence* etc. in Kelly 2011, pp.201ff) seem to grant such a *linear* interaction. This the more, as in the following model the weights can be varied from case to case (one might argue that linear weighting does *not* allow for a *synergy effect* of increasing one's degrees of belief due to consilience with that of one's peers, whereas geometric weighting does; however, Kelly, who argues for such an effect, does so only in the context of *peer agreement* and not *peer disagreement*):

Total Evidence View:

Among peers one's own position might be partly or fully influenced by the other peers' positions: 1. of (EPD) and

$$2. w_1^t + \dots + w_n^t = 1, \quad (\text{TEV})$$

$$w_j^t \geq 0 \quad (1 \leq j \leq n)$$

So, (TEV) just guarantees that the weights used in resolving epistemic peer disagreement allow for linear weighting.

The arguments for and against each of these views are well discussed in the debate on epistemic peer disagreement. We will not present and discuss them here in detail. Rather, given our formal model, we want to add a further argument to the debate which strikes us as decisive with respect to (EWV). We will then recap some arguments against (EWV) when we consider possible objections to our argument in section 5.

3 Inferential Skills As Reliabilities

One argument for the equal weight view (EWV) originates from indifference-considerations: If the epistemic attitudes of some peers are indistinguishable

with respect to their underlying evidence as well as their inferential skills, why should there be a difference with respect to their epistemic impact in updating one’s degrees of belief, once one becomes aware of a disagreement? The assumption that all agents share the same evidence is already *hard-wired* in our model: At each round all agents update on the same first order evidence. But how can we express that the peers have the same inferential skills? We want to suggest implementing this into the model by help of a reliabilistic measure of “inferential” or predictive success. The idea is as follows: Each agent has to make a prediction about the truth value of a proposition of a round t : p^t . These predictions $Pr_i^t(p^t)$ are based on the shared first order evidence as well as the individual inference or prediction method of the agent (Pr_i). We assume that afterwards all predictions are revealed to all agents and might serve as higher order evidence e^t for the same round t . Then each agent has to make again a prediction about p^t . Now, we assume that at the end of a round t the truth value of p^t is settled – and for [1548] simplicity of expression we assume also that it is revealed to all agents. So, the cases to which the argument presented here applies are not cases of deep disagreement; a characteristic of deep disagreement is that it cannot be resolved by help of further evidence. However, as we will indicate in the next section, there is a way of relaxing the assumption that the outcomes are revealed to the agents.

In this dynamics each round t consists of three phases: a phase of updating on first order evidence, a phase of updating on higher order evidence e^t , and a phase where the truth value is revealed. Figure 1 illustrates how the specified model of making inferences and predictions looks like.

t:	0	1	2	3	4	...
phases:	priors	Pr^1 e^1 p^1	Pr^2 e^2 p^2	Pr^3 e^3 p^3	Pr^4 e^4 p^4	...

Figure 1: Our model of peer disagreement: At each round all agents receive first order evidence and have to make their inferences based on this evidence: $Pr^t(p^t)$. Afterwards all agents receive information about the other agents’ inferences, i.e. higher order evidence e^t . They must make further inferences about the propositions in question which they might base on this further evidence. At the end of each round the truth value of the proposition(s) (p^t) in question is revealed.

Now, as we indicated above, given the truth values of the propositions in question one can measure the reliability of an agent’s predictions and inferences by tracking the average closeness of the agent’s predictions and inferences to the truth. Such a measure for *verifying forecasts* was already very early and quite prominently suggested by Brier (1950) in form of summing up the squared differences between the predicted value and the outcome as a score. In this paper we use the quadratic scoring measure because of the popularity

of this measure, but for technical reasons we use it in the exponent – for details on using different scoring measures see (Feldbacher-Escamilla [under revision](#), sect.3.4). Our model can be shown to be very robust regarding the exact choice of such a scoring measure. So, we define the reliability or success s_i^t of an inference or prediction method of agent i up to round t as follows (the truth value of the proposition p in question is given via $val(p)$, where truth is represented by $val(p) = 1$ and falsity by $val(p) = 0$):

$$s_i^t = \exp \left(\sum_{0 < u \leq t} 1 - (val(p^u) - Pr_i^u(p^u))^2 \right) \quad (1)$$

Note that in this equation one can find all elements of the three phases for some round t : The assessment of the agents given first order evidence in the first phase ($Pr_i^t(p^t)$), the higher order evidence of the second phase (which is $Pr_i^t(p^t)$ and the past reliability rate s_i^{t-1}), and the revealed truth value of the third phase ($val(p^t)$). An agent who gets all inferences absolutely right has maximum reliability of $\exp(t)$, whereas an agent who gets all inferences absolutely wrong has a reliability of 0. All other kinds of inferences are strictly in-between this interval.

Now, let us explicate the notion of an *epistemic peer* in this model. As we discussed in the preceding sections, two agents are epistemic peers iff they possess the same evidence and equal inference skills regarding the evidence. Since evidence sharing is [1549] hard-wired in the model, all agents of the model are peers in this respect. But how about the other relevant attribute? It seems plausible to assume that equal inferential skills can be expressed by equal reliabilities: According to this model, two agents are equally skilled regarding inferences on the basis of shared evidence at round t , if their reliabilities s_1^t and s_2^t match, i.e.: $s_1^t = s_2^t$. So, the question of how to update one’s degrees of belief on higher order evidence about one’s epistemic peers’ epistemic attitudes results in the question of how to update, given $s_i^t = s_j^t$ (for all $1 \leq i, j \leq n$). Now, we assume this condition for our model of epistemic peer disagreement by relativising (EPD) to cases of disagreement with peers that have the same reliabilities:

$$(EPD) \text{ holds for cases with } s_1^t = \dots = s_n^t$$

Note that the problem of epistemic peer disagreement as defined above is due to this presupposition relevant only for cases where the reliabilities s_i^t of the peers match. So, we explicate “same inferential skills” as having identical reliabilities. One might think that this is an assumption too strong in order to be put forward for the notion of *epistemic peers*, because in fact reliabilities will never match exactly. However, first of all, we think that once one is able to figure out some kind of track record, then different reliabilities also indicate different inferential skills and hence disqualify the agents in question as peers. Elga (2007), e.g., also notes that swamping arguments against the equal weight view are not that pressing once one takes into account that cases of peer disagreement show up less frequently than one might think. Secondly, to have equal

reliabilities does not imply that the individual inference methods are similar. Two inference methods can easily become equally reliable, although they constantly come up with different predictions; this happens, e.g., if one of them has a bias towards overestimation whereas the other has such a bias towards underestimation; or they just might get things right and wrong at different occasions. So, equal reliabilities of peers still allow for lots of variation among the peers. Thirdly, in applying the model one can always coarse-grain categories (e.g. the set of possible values which can be predicted or by clustering cases and taking their averages) and by this achieve more agreement in reliability rates – however, e.g., in coarse-graining categories the optimality results are also more restricted; for details cf. (Feldbacher-Escamilla 2017). And finally, although the equal weight view as described here covers, strictly speaking, only cases of disagreement among agents with identical reliabilities, a more general approach of equal weighting for peer and non-peer disagreement (as is outlined in Elga 2007) can be also shown to be covered by our optimality argument. We will say more on this in the section on possible objections. For now we ask the reader to consider the two explicated conditions for peerhood – namely update on shared first order evidence and identical reliabilities s^t – as clear-cut cases of peer disagreement.

4 The Optimality of Equal Weighting and The Suboptimality of Non-Equal Weighting

We want to show now that the equal weight view (EWV) is a specific instance of a general rule on incorporating higher order evidence which is proven to be optimal. [1550] We also want to show that the remain steadfast view (RSV) as well as the total evidence view (TEV) – in case it deviates from (EWV) – are, in terms of reliability, epistemically suboptimal. Here are the details: The reliability measure as defined in equation (1) can be considered as measuring the epistemic performance of an agent. As we indicated above, the best performance possible up to round t is given if an agent i 's reliability s_i^t is maximal, i.e. $\exp(t)$. This means that all of her inferences were correct up to t . For better comparison it is convenient to consider not only the *absolute* reliability of an inference or prediction method Pr_i up to [1551] round t , i.e. s_i^t , but also the *reliability rate*, which is simply a normalised average of the absolute reliabilities of all considered rounds: $s_i^t / \exp(t)$. Clearly the reliability rate $s_i^t / \exp(t)$ is within $[0, 1]$.

Now, a constraint which is often put forward for rational agents is *optimality*: An inference method Pr_i is epistemically rational, if its reliability rate is *long run optimal* compared to all available inference methods Pr_1, \dots, Pr_n in the sense that $s_i^t / \exp(t) \geq s_j^t / \exp(t)$ ($1 \leq j \leq n$), if t goes to infinity. Given this epistemic constraint, one can show that (EWV) satisfies it, whereas (RSV) as well as (TEV) fail to do so. For the case of (EWV) we will show this by employing a general optimality result of the so-called *theory of meta-induction*. For

showing the *suboptimality* of (RSV) and (TEV) we will provide an example where both of them fail to produce optimal inferences.

Let us start with the optimality of (EWV)! The theory of *meta-induction* was introduced by Schurz (2008) in order to address Hume’s problem of induction which poses a serious threat for the scientific task of making predictions and inferences. In order to address this problem, the theory of meta-induction provides a justification for inductive methods by stressing their ability to catch up with any inference method whatsoever in the long run. This meta-inductive justification of induction is twofold: First, a specific meta-method for selecting predictions of any accessible method in a reliability-based way is proven to be optimal in the sense described above. In a second step, taking for granted the past success of classical inductive methods, it follows that also selecting these methods for predictions of future events allows for optimal predictive success. This holds as long as there is no alternative method accessible which outperforms classical inductive methods.

Now, what is relevant for our investigation is not exactly the meta-inductive justification of inductive methods, but mainly the optimality of a meta-inductive method. For this reason we present the definition of such a method and prove its optimality in the appendix. Here is how it works (for comparison cf. the method EAW in Schurz 2019, the optimality results there are proven by help of a much more demanding mathematical apparatus): Let us think of the probability functions Pr_1, \dots, Pr_n as object inference methods in the sense that whenever they conditionalise on first order evidence, their inferences are functionally independent, i.e. the definition of such an inference method Pr_i based on first order evidence contains no reference to one of the other Pr_1, \dots, Pr_n . As described above (equation 1), for each such method Pr_i we can define a reliability measure s_i^t with a reliability rate $s_i^t / \exp(t)$. Now, the idea is to define a meta-method Pr_m whose inferences or predictions are weighted averages of the object-methods’ inferences and predictions, where the weights are based on a reliability measure as defined in equation (1). Given the reliability of a method Pr_i up to round t , i.e. given s_i^t , we get weights for the meta-method Pr_m by simple normalisation (it does not matter whether we take the absolute reliabilities or the reliability rate here):

$$w_i^{t+1} = \frac{s_i^t}{\sum_{j=1}^n s_j^t} \quad (2)$$

So, simply the more reliable an inference method is, the more weight it gains. Note that the weights for round $t + 1$ depend on the reliabilities up to round t .

Finally, by linear weighting according to these weights we can define the meta-method Pr_m as follows:

$$Pr_m^t(p^t) = \sum_{i=1}^n w_i^t \cdot Pr_i^t(p^t) \quad (3)$$

The weights for the meta-inductive inference or prediction Pr_m at round 0 might be arbitrarily chosen – it might be, e.g., the unweighted average of the object-level inferences.

Now, by proving main results of the theory of meta-induction and by ingeniously transferring important theorems of online learning theory to the theory of meta-induction, Schurz (2008) was able to show the following long run optimality result for such a meta-method – here $s_m^t / \exp(t)$ is the success measure or reliability rate of the meta-inductive method Pr_m :

Theorem 1. *In the long run, the reliability rate of the meta-inductive method approaches or outperforms that of the best available inference method(s).*

$$\lim_{t \rightarrow \infty} \left(\max \left(\frac{s_1^t}{\exp(t)}, \dots, \frac{s_n^t}{\exp(t)} \right) - \frac{s_m^t}{\exp(t)} \right) \leq 0$$

The theorem states that in the limit, i.e. in the long run, the difference between the reliability rate of the meta-inductive method and that of the most accurate method vanishes, or that the meta-inductive method even outperforms the most accurate method. In the following we provide an example illustrating the idea behind this theorem. In the appendix we provide an elementary proof of this theorem which is intended to be better accessible than most of the proofs of such results in this area of research.

The main “cause” relevant for success in terms of an outperforming reliability rate is that the weights for cooking up a prediction or inferring an estimation are based on the reliability rates. Perhaps this can be illustrated best by considering the case of a best competing inference method in the setting. This means that there is some point in time (i.e. some round) where the reliability rate of that method is no longer outperformed by any other method. The mechanism of reliability-based weighting is relatively transparent in this case: By considering the difference of the reliabilities in the exponent, every advantage – no matter how minimal it is – will increase the weight the best inference method gains until it has (or approaches) full weight. Figure 2 depicts such a case.

By theorem 1 we know that Pr_m is an inference method that performs optimally in the long run compared to all available inference methods Pr_1, \dots, Pr_n . [1552] Given our epistemic constraint, this provides a reason for considering Pr_m to be a rational inference method, i.e. to be epistemically justified.

Now, it is easy to see that Pr_m is an inference method based on higher order evidence only. Furthermore, Pr_m is a long run optimal method for all cases, cases of no disagreement, cases of disagreement, cases of disagreement among epistemic peers, and cases of disagreement among non-peers. What is important to note is that in the specific case of epistemic peer disagreement, Pr_m coincides with (EWV): Given $s_1^t = \dots = s_n^t$ as we presupposed for (EPD), it follows from the definition of Pr_m via equation (2) that $w_1^t = \dots = w_n^t = 1/n$. Hence, (EWV) instantiates Pr_m for the specific case of epistemic peer disagreement. Since Pr_m is shown to be optimal regarding all cases of agreement and




Figure 2: Example of reliability-based weighting; Pr_1 and Pr_2 are object inference methods; Pr_m is the meta inference method as defined above. At the beginning the inferences of Pr_2 are more accurate (until round 2); hence, at the beginning Pr_m puts more weight on Pr_2 's predictions. Little by little Pr_m assigns higher weight to the more and more accurate Pr_1 until it fully disregards the less reliable inference method Pr_2 (round 20). Regarding the diagrams: top: predicted values of the methods (the true value is supposed to be constantly 1: $val(p^t) = 1$); middle: reliability rates showing that the rate of inaccurate Pr_2 decreases rapidly, whereas that of Pr_1 approaches a stable rate (and will increase later on); bottom: reliability dependent weights of the methods, showing that the best method in the setting Pr_1 gains relatively soon full weight.

disagreement, (EWV) concerns the particular case of responding optimally to cases of epistemic peer disagreement. Hence, according to our epistemic constraint, (EWV) is epistemically justified regarding cases of epistemic peer disagreement.

It is important to note that here we claim only that (EWV) instantiates the meta-inductive Pr_m in cases of peer disagreement. We want to consider (EWV) here to be silent with respect to all other cases of dis-/agreement. So, this version of (EWV) is undetermined [1553] regarding cases where $s_i^t \neq s_j^t$. Perhaps the following externalist interpretation of our model helps to distinguish (EWV) better from the meta-inductive weighting strategy Pr_m , and hence identify our model of applying (EWV) as a model of dealing with peer disagreement (the following interpretation is intended for illustrative purposes; it is not intended for committing us to externalism): Assume that the reliabilities s_i^t, s_j^t are only accessible to an *external subject*, but not to the individual agents of the setting. Suppose furthermore that the individual agents get from the external subject only the information that they are peers, so, only the external subject knows that $s_i^t = s_j^t$, whereas this is not internal or transparent to the individual agents, because they did not keep track of reliabilities at all. Given this interpretation, an individual agent applying (EWV) is not at all like the

meta-inductive Pr_m , since the informational basis of the former consists only of knowledge about the assessment of the other agents and that they are peers, whereas the informational basis of the latter requires equation (1). Still, from the externalist subject’s perspective, the individual applying (EWV) is optimal with respect to cases of peer disagreement.

Let us come to a suboptimality proof of the alternative approaches to epistemic peer disagreement: (RSV) and (TEV). For this purpose it suffices to provide examples where these approaches fail to be optimal in the long run. It is easy to construct an “environment” which favours the competitors of (RSV) and (TEV), although both methods might, from time to time, catch up in reliability terms. Consider, e.g. a setting with two agents, having inference methods Pr_1 and Pr_2 , where the latter represents a *steadfaster* (RSV) or an agent considering the total evidence (TEV) which does not coincide with (EWV). Now, let us assume that out of three pairs of inference rounds, one is a round with epistemic peer disagreement and the other two are rounds with disagreement among non-peers. In general, for constructing an example with suboptimal behaviour we need to assume that there are cases of non-peers, because otherwise all reliabilities would be always equal and hence there would be no suboptimality. For providing a counterexample to (RSV) and (TEV) the important task is to prove suboptimality *due* to not equal weighting in cases of peer disagreement. And we can do so by thinking of inferences with reliabilities distributed as shown in table 1: At round u (and $u + 3, u + 6, \dots$) Pr_1 and Pr_2 are non-peers and Pr_2 , the steadfaster or total evidence strategist, outperforms Pr_1 ; at $u + 1$ (and $u + 4, u + 7, \dots$) both are peers and equally reliably; and at $u + 2$ (and $u + 5, u + 8, \dots$) they become non-peers again, but this time with Pr_1 outperforming Pr_2 . By averaging the reliabilities, one can see that Pr_1 outperforms Pr_2 , since in the long run (on average) it holds $s_1^{t \rightarrow \infty} = \exp(0.50) > \exp(0.497) = s_2^{t \rightarrow \infty}$. Again, note that the example is constructed in such a way that suboptimality results from not equally weighting in case of a peer disagreement: Agent 2 outperforms agent 1 before agent 1 becomes a peer and agent 2 performs suboptimally when deciding to not equally weight higher order evidence of agent 1’s epistemic attitude. Clearly, there are favourable settings for steadfasters and total evidence strategists too: remaining steadfast might result in a reliability series of $\exp(0.51), \exp(0.50), \exp(0.50)$ (instead of $\exp(0.51), \exp(0.50), \exp(0.48)$), but to prove the possibility of suboptimality in one situation is not in contrast with proving optimality in another one. The main problem is that there is a possibility of suboptimality for steadfasters and total evidence strategists *due* to not equal weighting in case of disagreement among peers.

[1554] The optimality result for (EWV) – as cited in theorem 1 – shows that such a case cannot appear if one performs equal weighting. Clearly, an agent performing (EWV) can also be suboptimal compared to her competitors. The simplest case one might think of is a setting where no peer disagreement shows up, because the agents’ reliabilities never match. However, this suboptimality of (EWV) is due to her being suboptimal in cases of non epistemic peer disagreement. Regarding cases of epistemic peer disagreement – which are the

	u	$u + 1$	$u + 2$	\dots
$val(p^t)$	1.0	1.0	1.0	\dots
$Pr_1^t(p^t)$	$1 - \sqrt{1 - .5}$	$1 - \sqrt{1 - .5}$	$1 - \sqrt{1 - .5}$	\dots
$Pr_2^t(p^t)$	$1 - \sqrt{1 - .51}$	$1 - \sqrt{1 - .5}$	$1 - \sqrt{1 - .48}$	\dots
s_1^t	$exp(.50)$	$exp(.50)$	$exp(.50)$	\dots
s_2^t	$exp(.51)$	$exp(.50)$	$exp(.48)$	\dots
			$\nwarrow \nearrow$ <i>peer disagreement</i>	

Table 1: Example of the suboptimality of (RSV) and (TEV) due to not weighting equally among one’s epistemic peers in case of epistemic peer disagreement: Pr_1 gets the inferences in 50% of the cases right, whereas Pr_2 is sometimes slightly better, then Pr_1 catches up and then, in the case of a peer disagreement, strategy Pr_2 of remaining steadfast or incorporating total evidence loses. Peer disagreement consists in equal reliabilities in round $u + 1$ and different predictions in round $u + 2$ (the relevant parameters are marked grey).

cases the debate is about – (EWV) is guaranteed to be optimal.

5 Possible Objections, Replies, and Restrictions

In this section we discuss possible objections to particular assumptions of our model, namely that the truth is always revealed, that one can employ a big enough track record, and that we consider peers via exactly matching reliabilities. Finally, we embed our argument in favour of the equal weight view (EWV) a little bit more into the traditional debate by showing how one can understand and address the already mentioned *swamping problem* within this approach.

Regarding deep disagreement and revealing the truth: One might object that in order to measure the reliabilities of the epistemic agents, our model presupposes that the truth values of the propositions are revealed at some point in time to the epistemic agents (via $val(p)$ in phase 3). However, in case of *deep disagreement* there might be no possibility to get to know these values. So, how can our model be employed in these cases? Although we think that one can address this problem in the particular applications due to the high degree of flexibility in interpreting the model, we also want to mention that a possible route consists in using an expanded meta-inductive framework which can cope also with so-called *intermittent prediction games* where the prediction series do not have to be complete. One can also devise a reliability measure for such games by relativising the reliability rate to the number of rounds where predictions and outcomes are available. The meta-inductive optimality result [1555] also transfers to these games (cf. Schurz 2019, sect.7.2). Now, if we interpret cases

of *deep* disagreement – i.e. cases in which the disagreement is never resolved, where the true outcome is not available – as such intermittent rounds, then we can also apply the model to cases of deep disagreement: The reliability rates are calculated by help of non-deep cases of disagreement and then applied to cases of deep disagreement too. And the optimality result states that equally weighting equally reliable inference methods allows for best predictive success in the long run. Hence, (EWV) is an optimal response in cases of deep disagreement too.

Regarding getting a big enough track record: Here one might wonder whether a big enough basis for calculating relevant reliability rates is always accessible (the model presupposes an arbitrary growing n). Clearly, in many important cases of peer disagreement it will be very hard to find such a big enough basis. So, e.g., in estimating the reliability of two disagreeing, let us say, physicians, one typically lacks track record of them regarding exactly similar cases. However, note that in our model the events underlying the prediction series need not be of the same type or in any other way closely connected to one another. The only thing that matters is that one considers the series of events (may they be of a very different type or not) to be of some relevance. Whatever this exactly means is left open by the framework; it is up to the specific application to cluster events into long enough series.

Regarding peerhood and exactly matching reliabilities: Now, as mentioned above, to restrict cases of peer disagreement to cases of exactly matching reliabilities is for many cases too restrictive and idealised. Above we have already outlined some possibilities to increase chances of matching reliabilities, namely coarse-graining of categories, clustering and averaging of predictions. These methods are only about an application of the model. There is, however, another possibility to de-idealise the model which is directly about the interpretation of the model: There is a straightforward expansion of the equal weight view which suggests to give more weight to more reliable agents and less weight to less reliable ones. Elga (2007, p.489) describes the expansion as follows:

“At the start, the equal weight view applied only to cases in which you initially count your advisor as a peer – as equally likely to be right, on the supposition that the two of you end up disagreeing. But the modified view also applies to cases in which you initially count an advisor as an epistemic superior – as being more than 50% likely to be right, on the supposition that the two of you end up disagreeing. Likewise, the view applies to cases in which you initially count an advisor as an epistemic inferior.”

The idea is that according to an expanded version of (EWV) what matters is not that the weights for all individuals are *equal* ($w_1^t = \dots = w_n^t$), but that the weights and the (normalised) reliabilities are *equal*. This expansion covers the strict case of peer disagreement with $s_1^t = \dots = s_n^t$ as well as all other cases of disagreement. Now, by this one can also relax the second condition we put forward for *peerhoodness*, namely the condition of identical reliabilities. One might consider

peerhoodness as a notion which comes in degrees or which allows not only for strictly identical, but also similar inferential skills ($s_1^t \approx \dots \approx s_n^t$). Regardless of how exactly one relaxes this condition, if one considers the expanded version of (EWV) where equality is not only [1556] demanded for the weights in case of strictly matching reliabilities, but where equality is in general demanded for the weights and the respective (normalised) reliabilities, then our optimality argument favours also an expanded version of (EWV) with less ideal assumptions about peerhood.

Regarding swamping: Regarding the problem of higher order evidence *swamping* first order evidence our framework shows why and how higher order evidence can relevantly outrule first order evidence and might even swamp it: Inferences based on first order evidence are prone to being suboptimal, whereas inferences based on higher order evidence drastically decrease possibilities of being suboptimal relative to one's fellows. As we have shown, this holds particularly for cases of disagreement amongst peers. Regarding disagreement amongst non-peers, there is also a conclusion one can draw from our framework: The long run optimality result for Pr_m is based on the proof of an upper bound for differences in the reliability rate between the individual first order inferences and that one which is based on higher order evidence. As we show in the proof in the appendix, for Pr_m the following exact boundary holds with respect to any inference method Pr_i of Pr_1, \dots, Pr_n for some known t (step 11 of the proof):

$$\frac{s_i^t}{\exp(t)} - \frac{s_m^t}{\exp(t)} \leq \frac{\sqrt{2 \ln(n)} \cdot t}{e^{\sqrt{2 \ln(n)} \cdot t}}$$

In the limit this term goes to 0 which means that s_m approaches also the reliability of the best agents in the setting. However, for the short run, i.e. t not arbitrarily high, n , i.e. the number of agents in the setting, has a relevant influence on the performance. The more non-peers there are in the setting, the more the higher order strategy is prone to errors. So, in such cases first order evidence might easily outrule higher order evidence, i.e. in this cases higher order evidence should not swamp first order evidence. Notice, however, that this concerns cases of disagreement amongst non-peers. And for these cases (EWV) leaves the choice of weights for incorporating first and higher order evidence completely open, since, as we have presupposed for EPD, cases of peerhood were characterised as cases where the individual reliabilities s_1^t, \dots, s_n^t are identical or close to each other.

We have named here a couple of worries one might have with our approach to the problem of epistemic disagreement. Due to lack of space we can provide only hints on how these problems might be addressed: It seems that the problem of *deep* disagreement can be accounted for by including intermittent prediction games into the model where the truth is not always released. The problem of generating a big enough basis for calculating a reliabilistic track record might be addressed by clustering different types of events to single event series; for the problem of assuming matching reliabilities for peerhood we indi-

cated how an expansion of the view as suggested, e.g., by Elga (2007) covers also cases of approximative matching reliabilities or, more generally, reliability dependent weighting (where in the expanded version *equal in equal weight view* is about equating the weights with the reliabilities). And for the swamping problem our suggestion is that in case of a small record first order evidence should not be swamped, but that with an increasing amount of such a record higher order evidence becomes more influential due to its feature of guaranteeing optimality. [1557]

6 Conclusion

In the debate of epistemic peer disagreement the central question concerns the problem of how to incorporate or weight higher order evidence about epistemic peers disagreeing with one's epistemic attitudes. In this paper we have presented a model for such disagreement that frames this problem as a problem of updating one's credences given such higher order evidence. We have identified the conditions for epistemic peers in updating on one and the same set of first order evidence, and in having the same inferential skills regarding this evidence, measured by a reliability track record.

We were able to define in this model the three traditional approaches: the remain steadfast view (*RSV*) which suggests to ignore such evidence in updating by assigning it zero weight. The equal weight view (*EWV*) which suggests to update on such evidence by equally weighting it. And finally, the total evidence view (*TEV*) which suggests to consider higher order evidence simply as just another kind of evidence.

Since updating on higher order evidence is a social epistemological action, we have put forward an optimality constraint for such actions by aiming at optimality: One's update on higher order evidence should be such that (in the long run) one's inferences are optimal in the sense that they are the most reliable ones compared to the other inference methods of the setting. By employing the framework of *meta-induction* from Schurz (2008) we were able to show that (*EWV*) satisfies this optimality constraint for cases of epistemic peer disagreement, whereas (*RSV*) and (*TEV*) fail to do so. This adds a new argument to the debate on epistemic peer disagreement which seems to be decisive with respect to (*EWV*).

Acknowledgements

I would like to thank Peter Brössel for collaborating on this topic, his support in establishing the main ideas of the model, and his valuable feedback on first drafts of the paper. My thanks go also to Anna-Maria Eder, Thomas Grundmann, Christoph Jäger, Gerhard Schurz, and two anonymous referees for important discussions and helpful comments.

Appendix

Here is a proof of theorem 1: The main strategy of the proof is to apply inequalities such that the differences of the reliabilities are narrowly bounded. As we demonstrate now, reliability-based weighting allows for an optimal bound in the sense that in the limit such weighting cannot be outperformed by any other inference method in terms of the reliability rate. As we have argued above, the equal weight view is such a reliability-based weighting method; for this reason the proven guarantee holds also for the equal weight view. Here is a quite explicit sketch of the proof:

Proof. We aim at proving that reliability-based weighting allows for optimality in the sense that the reliability rate of such a predictor is not outperformed by that of its competitors. In order to do so, what one typically does in online learning theory is to characterise the difference between the competing predictors and that of the reliability-based predictor by help of a learning parameter η which is a function of the number of rounds t , and which grows sublinearly with t . If such a characterisation succeeds, then the difference of the reliability rate, which is the difference of the reliabilities divided [1558] by t , also grows sublinearly only and vanishes in the limit; this means that by help of such a characterisation the reliability-based predictor is shown to be not outperformed by any other predictor in the limit. As it turns out, one can characterise such differences in reliabilities by help of choosing $\eta = \sqrt{\frac{2 \cdot \ln(n)}{T}}$. Here T is an arbitrary round and sometimes also called the *prediction horizon* up to which a boundary is proven. In order to generalise this boundary to any round t , one needs, in a second step, to get rid of the exact choice of T by employing the so-called *doubling trick*, according to which for each round t it is assumed that the prediction horizon T doubles; this assumption increases the bound a bit, but does not change anything regarding the limiting case, and hence allows for proving a general optimality result too. In the following proof we demonstrate the first part (for arbitrary T); the second part of applying the doubling trick can be recapitulated by help of (Mohri, Rostamizadeh, and Talwalkar 2012, p.158).

1. Define $l_i^t = l(\text{val}(p^t), \text{Pr}_i^t(p^t|e^t)) = (\text{val}(p^t) - \text{Pr}_i^t(p^t|e^t))^2$ as the quadratic loss of predictor i at round t with respect to predicting p^t (on the basis of first order evidence e^t).
2. Let us assume that \exp is the basis e^t , where η is a learning parameter defined as $\eta = \sqrt{\frac{2 \cdot \ln(n)}{T}}$ and T is some fixed *prediction horizon*; here we aim only at proving a bound up to an arbitrarily high prediction round T (by help of the mentioned *doubling trick* – where one considers at each round as prediction horizon two times the round number – one can get rid of this parameter). Furthermore let l be convex (the squared difference function we used in our model is convex – more on convexity see below, when we use this property).

3. According to equation (1), the reliability s_i^t of a prediction method Pr_i up to t is defined as the cumulative non-averaged score of Pr_i up to t in the exponent. We can reformulate this measure recursively as follows: $s_i^{t+1} = s_i^t \cdot \exp(1 - l(\text{val}(p^t), Pr_i^t(p^t|e^t)))$ (keep in mind that below we will take \exp according to 2. to be e^η). By this re-formulation we get the following equalities about the ratio of the denominators used in normalisation (the normalising denominator for t and that of $t - 1$) in the definition of the weights according to equation (2) – we write for short l_i^t for $l(\text{val}(p^t), Pr_i^t(p^t|e^t))$:

$$\begin{aligned} \frac{\sum_{i=1}^n s_i^t}{\sum_{j=1}^n s_j^{t-1}} &= \frac{\sum_{i=1}^n s_i^t}{\sum_{j=1}^n s_j^{t-1}} = \frac{\sum_{i=1}^n \frac{s_i^{t-1} \cdot e^{-\eta \cdot l_i^t} \cdot e^\eta}{\sum_{j=1}^n s_j^{t-1}}}{\sum_{j=1}^n s_j^{t-1}} \\ &= \sum_{i=1}^n w_i^t \cdot e^{-\eta l_i^t} \cdot e^\eta \end{aligned}$$

4. By the inequality $e^{-x} \leq 1 - x + \frac{x^2}{2}$ (valid for all $x \geq 0$) we get the instance:

$$e^{-\eta \cdot l_i^t} \leq 1 - \eta \cdot l_i^t + \frac{\eta^2 \cdot l_i^{t2}}{2}$$

5. [1559] By substituting the right term in the inequality of 4. for the respective e -term in 3. we get:

$$\frac{\sum_{i=1}^n s_i^t}{\sum_{j=1}^n s_j^{t-1}} \leq \sum_{i=1}^n w_i^t \cdot \left(1 - \eta \cdot l_i^t + \frac{\eta^2 \cdot l_i^{t2}}{2} \right) \cdot e^\eta$$

and by arithmetic transformation:

$$\leq \left(\sum_{i=1}^n w_i^t - \left(\eta \cdot \sum_{i=1}^n (w_i^t \cdot l_i^t) - \frac{\eta^2}{2} \cdot \sum_{i=1}^n (w_i^t \cdot l_i^{t2}) \right) \right) \cdot e^\eta$$

By the normalisation of w : $\sum_{i=1}^n w_i^t = 1$, so:

$$\leq \left(1 - \left(\eta \cdot \sum_{i=1}^n (w_i^t \cdot l_i^t) - \frac{\eta^2}{2} \cdot \sum_{i=1}^n (w_i^t \cdot l_i^{t2}) \right) \right) \cdot e^\eta$$

By taking the \ln on both sides of the inequality:

$$\ln \left(\frac{\sum_{i=1}^n s_i^t}{\sum_{j=1}^n s_j^{t-1}} \right) \leq \ln \left(1 - \left(\eta \cdot \sum_{i=1}^n (w_i^t \cdot l_i^t) - \frac{\eta^2}{2} \cdot \sum_{i=1}^n (w_i^t \cdot l_i^{t2}) \right) \right) + \eta$$

6. By the inequality $e^{-x} \geq 1 - x$ (valid for any x) we get $\ln(e^{-x}) \geq \ln(1 - x)$ and hence $-x \geq \ln(1 - x)$. So, as an instance:

$$-\left(\eta \cdot \sum_{i=1}^n (w_i^t \cdot l_i^t) - \frac{\eta^2}{2} \cdot \sum_{i=1}^n (w_i^t \cdot l_i^{t2})\right) \geq \ln\left(1 - \left(\eta \cdot \sum_{i=1}^n (w_i^t \cdot l_i^t) - \frac{\eta^2}{2} \cdot \sum_{i=1}^n (w_i^t \cdot l_i^{t2})\right)\right)$$

7. By substituting the left (upper) term in the inequality of 6. for the right term in the inequality in 5. we get: [1560]

$$\ln\left(\frac{\sum_{i=1}^n s_i^t}{\sum_{j=1}^n s_j^{t-1}}\right) \leq -\left(\eta \cdot \sum_{i=1}^n (w_i^t \cdot l_i^t) - \frac{\eta^2}{2} \cdot \sum_{i=1}^n (w_i^t \cdot l_i^{t2})\right) + \eta$$

and by arithmetic transformation:

$$\leq \frac{\eta^2}{2} \cdot \underbrace{\sum_{i=1}^n (w_i^t \cdot l_i^{t2})}_{\leq 1} - \eta \cdot \sum_{i=1}^n (w_i^t \cdot l_i^t) + \eta$$

... due to $\sum_{i=1}^n w_i^t = 1$, and $l \in [0, 1]$, so:

$$\leq \frac{\eta^2}{2} \cdot 1 - \eta \cdot \sum_{i=1}^n (w_i^t \cdot l_i^t) + \eta$$

8. So, we arrived at the inequality (from 7.):

$$\ln\left(\sum_{i=1}^n s_i^t\right) - \ln\left(\sum_{j=1}^n s_j^{t-1}\right) \leq \frac{\eta^2}{2} - \eta \cdot \sum_{i=1}^n (w_i^t \cdot l_i^t) + \eta$$

Now we can sum up each side of the inequality from 1 to T :

$$\sum_{t=1}^T \left(\underbrace{\ln\left(\sum_{i=1}^n s_i^t\right)}_{=def C_t} - \underbrace{\ln\left(\sum_{j=1}^n s_j^{t-1}\right)}_{=def C_{t-1}} \right) \leq \sum_{t=1}^T \left(\frac{\eta^2}{2} - \eta \cdot \sum_{i=1}^n (w_i^t \cdot l_i^t) + \eta \right)$$

$$= (C_T - C_{T-1}) + \dots + (C_2 - C_1) + (C_1 - C_0)$$

$$=_{C_T - C_0} \frac{T \cdot \eta^2}{2} - \eta \cdot \sum_{t=1}^T \sum_{i=1}^n (w_i^t \cdot l_i^t)$$

So, we arrive at:

$$\ln\left(\sum_{i=1}^n s_i^T\right) - \underbrace{\ln\left(\sum_{j=1}^n s_j^0\right)}_{=n} \leq \frac{T \cdot \eta^2}{2} - \eta \cdot \sum_{t=1}^T \sum_{i=1}^n (w_i^t \cdot l_i^t) + T \cdot \eta$$

Hence:

$$\ln \left(\sum_{i=1}^n s_i^T \right) - \ln(n) \leq \frac{T \cdot \eta^2}{2} - \eta \cdot \sum_{t=1}^T \sum_{i=1}^n (w_i^t \cdot l_i^t) + T \cdot \eta$$

[1561] Recall, s_i^t is the cumulative score up to t in the exponent and we are after the bound for the difference with respect to the best predictor, hence we concentrate on the predictor with maximal cumulative score up to T : Let us denote this predictor with ' b ' ($b = \text{argmax}_{i \in \{1, \dots, n\}} (\sum_{t=1}^T 1 - l_i^t)$). If there are more, then we can randomly pick one. From above we get:

$$\ln(s_b^T) - \ln(n) \leq \frac{T \cdot \eta^2}{2} - \eta \cdot \sum_{t=1}^T \sum_{i=1}^n (w_i^t \cdot l_i^t) + T \cdot \eta$$

9. By definition of s :

$$\ln(s_b^T) = \ln \left(\exp \left(\eta \cdot \sum_{t=1}^T (1 - l_b^t) \right) \right) = \eta \cdot \sum_{t=1}^T (1 - l_b^t)$$

By substituting the right term in the last inequality in 8. we get:

$$\eta \cdot \sum_{t=1}^T (1 - l_b^t) - \ln(n) \leq \frac{T \cdot \eta^2}{2} - \eta \cdot \sum_{t=1}^T \sum_{i=1}^n (w_i^t \cdot l_i^t) + T \cdot \eta$$

And by arithmetical transformation:

$$\sum_{t=1}^T \sum_{i=1}^n (w_i^t \cdot l_i^t) - \sum_{t=1}^T l_b^t \leq \frac{T \cdot \eta}{2} + \frac{\ln(n)}{\eta}$$

If we substitute for η in accordance with 2: $\eta = \sqrt{\frac{2 \cdot \ln(n)}{T}}$, we get:

$$\sum_{t=1}^T \sum_{i=1}^n (w_i^t \cdot l_i^t) - \sum_{t=1}^T l_b^t \leq \sqrt{2 \cdot \ln(n) \cdot T}$$

Now, what is left is to employ the grey marked term for proving a bound for the difference between the cumulative l_m and l_b (which is called *regret* in the machine learning literature).

10. According to equation (3), P_m predicts as follows: $P_m^t(p^t) = \sum_{i=1}^n w_i^t \cdot P_i^t(p^t)$. Hence its loss is: $l \left(\text{val}(p^t), \sum_{i=1}^n (w_i^t \cdot P_i^t(p^t)) \right)$. And hence its cumulative loss is:

$$\sum_{t=1}^T l \left(\text{val}(p^t), \sum_{i=1}^n (w_i^t \cdot P_i^t(p^t)) \right)$$

[1562] Since l is convex (according to 2.), we get:

$$l\left(\text{val}(p^t) \sum_{i=1}^n (w_i^t \cdot P_i^t(p^t))\right) \leq \sum_{i=1}^n (w_i^t \cdot l(\text{val}(p^t), P_i^t(p^t)))$$

(I.e.: The loss of a weighted average of predictions is smaller than or equal to the weighted average of the losses of the predictions.) Hence, from the last inequality in 9. and the convexity of l we get:

$$\underbrace{\sum_{t=1}^T \left(l\left(\text{val}(p^t), \sum_{i=1}^n (w_i^t \cdot P_i^t(p^t))\right) \right)}_{\text{regret of } P_m \text{ with respect to } P_b} - \sum_{t=1}^T l_b^t \leq \sqrt{2 \cdot \ln(n) \cdot T}$$

Since P_b was the method with maximal absolute reliability up to T , i.e. least cumulative loss up to T (we defined b this way in 8.), this regret bound holds also with respect to all other predictors.

11. Now we can expand the left part of this inequality by $+T - T$ and include these terms into sum-terms such that we get:

$$\underbrace{\sum_{t=1}^T (1 - l_b^t)}_{s_b^T} - \underbrace{\sum_{t=1}^T (1 - l_m^t)}_{s_m^T} \leq \sqrt{2 \cdot \ln(n) \cdot T}$$

This implies for the reliability rates:

$$\frac{s_b^T}{\exp(T)} - \frac{s_m^T}{\exp(T)} \leq \frac{\sqrt{2 \cdot \ln(n) \cdot T}}{\exp(T)} = \frac{\sqrt{2 \cdot \ln(n) \cdot T}}{e^{\sqrt{2 \cdot \ln(n) \cdot T}}}$$

12. Hence:

$$\lim_{t \rightarrow \infty} \left(\frac{s_b^t}{\exp(t)} - \frac{s_m^t}{\exp(t)} \right) \leq 0$$

□

[1563]

References

- Brier, Glenn W. (1950). "Verification of Forecasts Expressed in Terms of Probability". In: *Monthly Weather Review* 78.1, pp. 1–3.
- Brössel, Peter and Eder, Anna-Maria A. (2014). "How to Resolve Doxastic Disagreement". In: *Synthese* 191.11, pp. 2359–2381. DOI: [10.1007/s11229-014-0431-4](https://doi.org/10.1007/s11229-014-0431-4).

- Christensen, David (2007). “Epistemology of Disagreement: The Good News”. In: *Philosophical Review* 116.2, pp. 187–217. DOI: [10.1215/00318108-2006-035](https://doi.org/10.1215/00318108-2006-035).
- (2010). “Higher-Order Evidence”. English. In: *Philosophy and Phenomenological Research* 81.1, pp. 185–215. URL: <http://www.jstor.org/stable/20779554>.
- Dietrich, Franz and List, Christian (2016). “Probabilistic Opinion Pooling”. In: *The Oxford Handbook of Probability and Philosophy*. Ed. by Hájek, Alan and Hitchcock, Christopher. Oxford: Oxford University Press.
- Easwaran, Kenny, Fenton-Glynn, Luke, Hitchcock, Christopher, and Velasco, Joel D. (2016). “Updating on the Credences of Others: Disagreement, Agreement, and Synergy”. In: *Philosophers’ Imprint* 16.11, pp. 1–39.
- Elga, Adam (2007). “Reflection and Disagreement”. English. In: *Noûs* 41.3, pp. 478–502. DOI: [10.1111/j.1468-0068.2007.00656.x](https://doi.org/10.1111/j.1468-0068.2007.00656.x).
- Feldbacher-Escamilla, Christian J. (2017). “Optimisation in a Synchronised Prediction Setting”. In: *Journal for General Philosophy of Science* 48.3, pp. 419–437. DOI: [10.1007/s10838-017-9379-7](https://doi.org/10.1007/s10838-017-9379-7).
- (under revision). *Epistemic Engineering. Uncovering the Logic of Deceivability and Meta-Induction*. book manuscript.
- Feldman, Richard (2007). “Reasonable Religious Disagreements”. In: *Philosophers Without God. Mediation on Atheism and Secular Life*. Ed. by Antony, Louise. Oxford: Oxford University Press, pp. 194–214.
- Kelly, Thomas (2011). “Peer Disagreement and Higher Order Evidence”. In: *Social Epistemology. Essential Readings*. Ed. by Goldman, Alvin I. and Whitcomb, Dennis. Oxford: Oxford University Press, pp. 183–217.
- Lehrer, Keith and Wagner, Carl (1981). *Rational Consensus in Science and Society. A Philosophical and Mathematical Study*. Dordrecht: Reidel Publishing Company.
- Mohri, Mehryar, Rostamizadeh, Afshin, and Talwalkar, Ameet (2012). *Foundations of Machine Learning*. Cambridge, Massachusetts: The MIT Press.
- Nissan-Rozen, Ittay and Spectre, Levi (2017-12). “A Pragmatic Argument Against Equal Weighting”. In: *Synthese*. DOI: [10.1007/s11229-017-1651-1](https://doi.org/10.1007/s11229-017-1651-1).
- Rosen, Gideon (2001). “Nominalism, Naturalism, Epistemic Relativism”. In: *Noûs* 35, pp. 69–91. DOI: [10.1111/0029-4624.35.s15.4](https://doi.org/10.1111/0029-4624.35.s15.4).
- Schurz, Gerhard (2008). “The Meta-Inductivist’s Winning Strategy in the Prediction Game: A New Approach to Hume’s Problem”. In: *Philosophy of Science* 75.3, pp. 278–305. DOI: [10.1086/592550](https://doi.org/10.1086/592550).
- (2012). “Meta-Induction and the Problem of Fundamental Disagreement”. In: *Epistemology: Contexts, Values, Disagreement. Proceedings of the 34th International Ludwig Wittgenstein-Symposium in Kirchberg am Wechsel, Austria 2011*. Ed. by Jäger, Christoph and Löffler, Winfried. Frankfurt: Ontos, pp. 343–354.
- (2019). *Hume’s Problem Solved. The Optimality of Meta-Induction*. Cambridge, Massachusetts: The MIT Press.