

DISAGREEMENT AND A FUNCTION EQUAL WEIGHT VIEW

Christopher A. Vogel

Vogel, C.A. (2022) Disagreement and a Function Equal Weight View.
Disputatio, 14(65): pp. 157–94

(please cite original source: <https://doi.org/10.2478/disp-2022-0009>)

Abstract

If a colleague of mine, whose opinion I respect, disagrees with me about some claim, this might give me pause regarding my position on the matter. The Equal Weight view proposes that in such cases of peer disagreement I ought to give my colleague’s opinion as much weight as my own, and decrease my certainty in the disputed claim. One prominent criticism of the Equal Weight view is that treating higher-order (indirect) evidence in this way invariably swamps first-order (direct) evidence. While the opinions of our peers matter in our deliberations, the Equal Weight view counter-intuitively requires that evidence of mere disagreement is more important than standard kinds of evidence. I offer a proposal for how we should idealize epistemic agents that identifies the variable feature of disagreements that accounts for the shifting significance of direct and indirect evidence in different disagreement contexts. Specifically, by idealizing epistemic agents as deriving functions that characterize the non-subjective relationship between a body of evidence and the reasonableness of believing the various propositions supported by that evidence, we can accommodate the intuition to compromise that motivates the Equal Weight view, without accepting the counter-intuitive results.

Keywords

disagreement, Equal Weight view, indirect evidence, uniqueness

1 Introduction

Epistemology¹ takes the reasonableness of our beliefs, the feature of belief formation that renders a belief justified, as a primary object of study. A recent revival of the topic of disagreement brings this aspect of epistemology to center stage, considering whether disagreement itself is a kind of evidence that should bear on belief formation. The space of possible positions regarding the evidential value of disagreement is exhausted by the continuum between two positions: the Direct Evidence view, and the Indirect Evidence view.² The Direct Evidence view denies that the opinions of other agents (i.e., peers, gurus, etc.) bear on the justification of a belief. Whether one is justified in believing some (weighted) proposition depends solely on non-agential evidence. Such *direct* evidence determines what beliefs are justified. Put another way, the opinions of epistemic agents, on the Direct Evidence view, do not count as evidence. On the other extreme, the Indirect Evidence view indicates that one forms justified beliefs *only* by appealing to the judgments of one's epistemic colleagues. Direct evidence, on the Indirect Evidence view, has no bearing on the justification for a belief, apart from the role such evidence plays in developing indirect evidence.

Intuitions push us to reject both of these extremes. The opinions of (at least some of) our epistemic colleagues seemingly matter in determining whether we are justified in holding a belief. But completely ignoring our own assessment of the direct evidence seems disingenuous, and in at least some cases we are compelled to accommodate our peer. This intuition motivates the Equal Weight (EW) view, which tells us to give the indirect evidence weight in proportion to the perceived reliability of its source. When that source is a peer, we must give our peer's assessment equal weight to our own (Elga, 2007). Kelly (2010) highlights several counter-intuitive implications of the EW view, all stemming from the *swamping* character given to indirect evidence. While epistemic humility motivates the *intuition to compromise*,

¹I am indebted to Jason Christie, Dan Moller, and Eric Pacuit for helpful discussion about this project at various stages. My thanks also to two anonymous reviewers, one at this journal and one at another, for insightful and encouraging feedback.

²Elga (2010) labels these two positions as the Stubborn view and the Conciliatory view, respectively. Less contentiously, Christensen (2011) labels these two views Steadfast and Conciliatory. Maximally Steadfast views hold that indirect evidence carries no epistemic weight, while maximally Conciliatory views hold that indirect evidence carries all the relevant weight.

accepting the general ascription to modify one's beliefs in the face of *any* indirect evidence counter-intuitively eliminates the import of direct evidence. This reflects a contrary, *steadfast* intuition, whereby we, at least sometimes, seem justified in ignoring the indirect evidence.

The positive proposal Kelly (2010) offers in response to the EW view however fails to explain in general *why* we ought to consider indirect evidence in some contexts, yet not in others. On his Total Evidence view the relevance of direct and indirect evidence in justifying a belief varies from case to case, depending on the character of the total available evidence. Contextualizing the solution to peer disagreement at such a fine grain admits that there is no general account of how indirect evidence should be considered. This concedes too much. In this paper I offer such a general account of disagreement in keeping with the EW view insofar as it acknowledges the appeal of the intuition to compromise, but does so without sliding into the counterintuitive implications that would ignore direct evidence entirely.

The core move on the proposed Functional Equal Weight (FEW) view relocates the target of our compromising intuition. Since we routinely evaluate evidence as offering support for some beliefs and not others, we must thereby assess which beliefs are reasonable to adopt given our available evidence. These assessments, or representations, of the *reasonableness* of a belief given a body of evidence, I argue, are the appropriate targets of the intuition to compromise. The upshot to this view is that we can maintain the intuitive pull of the intuition to compromise even though the mere fact of disagreement often fails to supply sufficient information to aptly accommodate the opinions of peers in reassessing the reasonableness of our beliefs. This shift away from credences and toward reasonableness forestalls the swamping problem with the Equal Weight view, by explaining why we are justified in ignoring a peer's registered disagreement and remain steadfast in contexts where the direct evidence is complex.

I begin by explicitly outlining the assumptions that motivate the problem of disagreement, primarily those relating to the non-subjective relation evidence bears to the reasonableness of a belief. In doing so I will indicate the content of the positive proposal offered in this paper: that the intuition to compromise pertains to the reasonableness of our beliefs, not simply their credences. I then turn to the Equal Weight view espoused in Elga (2007) as a means of highlighting both the intuitions that compel us to adopt the view, and the intuitions that inform the troubling swamping consequences, as argued in Kelly (2010). I then develop the Functional Equal Weight view

that remedies these problems by abstractly characterizing epistemic agents as primarily concerned with deriving a characteristic function that describes the non-subjective relationship between a body of evidence, the weighted propositions that evidence supports, and the reasonableness of believing any of those propositions. I then show how the FEW view resolves the central challenge of disagreement, giving quarter to the relevance of the indirect evidence in aggregating reasonableness representations, without letting the indirect evidence swamp the direct evidence. What will emerge is that the mere fact of disagreement about a proposition serves as weak evidence for determining the characteristic function of the relevant reasonableness representations. I will argue that when lacking the information required to perform the mandated compromising operation over reasonableness representations, one is justified in remaining steadfast since no alternative means of incorporating the indirect evidence is available in the paradigm contexts that license the steadfast intuition. This explains why we are rarely compelled to adjust our views in the face of disagreement.

2 Disagreement, Assumptions, and the Proposal

2.1 Assumptions

Each of us serves as a locus of belief formation. We observe the world we occupy, and form beliefs based on the evidence we thereby collect. Noting the rather banal fact that we often err in this task by forming false beliefs, our fallibility gives rise to the problem of disagreement. If I accept that I am fallible in evaluating how some body of evidence pertains to the truth of a proposition, discovering that another evaluator disagrees with my conclusions renders salient the possibility that I have made a mistake. The problem of disagreement is the task of determining the (normative) epistemic force of this discovery. In some cases, overcoming this problem is a simple matter. If the disputant is ignorant, stupid, or perhaps intoxicated, I seem justified in ignoring his registered disagreement. However, if the disputant is someone whose opinion I revere, someone I consider as intelligent or informed as I am, dismissing their conclusions outright seems (at best) hasty. The fictional economist's monologue below seems both familiar and compelling:

I think that interest rates will move down. My colleague disagrees. Furthermore, my colleague has examined all the same information I have, knows as much as I do about the issue, and is as well trained as I am. This worries me. After all, I take his opinion seriously and would welcome learning that he agrees with my view of this matter. Is it reasonable for me to retain my belief in light of this disagreement? Or is some adjustment rationally required? (Feldman & Warfield, 2010, 2)

In effect, our economist believes her colleague is on *par* with herself regarding the relevant epistemic variables. Consider such on *par* cases as cases of epistemic *peer* disagreement.³ These cases pose the central problem of this paper: how should we respond to the indirect evidence offered by our epistemic peers when they disagree with us, given our own fallibility?

Our economist is worried that she has *erred* in assessing the relevant evidence. Importantly, the very notion of error presupposes that there is some non-subjective standard to which one must comport. That is, if the reasonableness of the economist's judgments about interest rate growth were merely beholden to her own subjective assessment of the evidence, it is hard to see how she could ever be justifiably worried about being unreasonable or mistaken, *modulo* concerns of incapacity (e.g. impairment, intoxication, etc.). Thus, by entertaining the thought that her colleague's opinion matters, the economist presupposes that the appropriateness of her belief is determined (somehow or other) by factors beyond her subjective assessment of the evidence.

Of course one could think that the standard by which the economist wishes to comport is the truth. Maybe what worries the economist is that she is wrong about what the future will bring. If it turns out that interest rates rise in the near future, she will have made a false proclamation, and that (we might think) is wherein her potential error (and worry) lies. Her colleague's registered disagreement seems like evidence that she might have made such a mistake. But evidence can be misleading, making truth in some cases hard to come by. The laws of Newtonian Mechanics, so we have learned, are (strictly

³There are various ways in which we might formulate the notion of an epistemic peer. Following Feldman (2006) we might consider an epistemic peer to be "another person, every bit as sensible, serious, and careful as oneself", or we might consider someone an epistemic peer if "conditional on our disagreeing [they] are equally likely to be mistaken [as we are]" (Elga, 2007, 9). What follows from these various formulations is unclear, but what I argue here appears compatible with these various accounts.

speaking) false. But Newton was no fool, and characterizing his beliefs about the motion of physical bodies as unreasonable would seem to miss the mark.⁴ Likewise, the epistemology (and practical rationality) literature is filled with thought-experiments of people forming true beliefs for what we consider *bad* reasons. If, in 2005, I came to believe that Donald Trump will become a US President because my five-year-old daughter asserts as much after watching reality television, I have erred as an epistemic agent. Presumably I have erred despite the accuracy of my daughter’s prediction. Put more succinctly, people can have unreasonable beliefs, even if those beliefs are true, and likewise have reasonable beliefs that are false. For these reasons I assume that the intuitions motivating the trouble with cases of peer disagreement presuppose that the reasonableness of a particular belief is determined by non-subjective factors, and that there is some non-subjective relationship between a body of evidence and what one can reasonably believe based on that evidence.⁵

2.2 Proposal

The disagreement (and formal epistemology) literature often idealizes epistemic agents in certain ways. The most obvious is in treating human beliefs as confidence-proposition pairs, of the Bayesian sort. If I strongly believe the Earth to be a spheroid, this belief is treated abstractly as a confidence level from 0–1 attached to a proposition P ‘that the Earth is a spheroid’—or more succinctly, my weighted belief is $.8(P)$. This abstraction is meant to approximate a genuine phenomenon, namely that our beliefs admit to degrees of strength. But one would be hard pressed to defend the view that actual human beliefs admit to the kind of precision implicit in such formal treatment.

⁴Christensen (2011) marks a distinction between *doxastic* rationality and *propositional* rationality. One is propositionally rational if they come to believe (somehow or other) the proposition most supported by the available evidence. One is doxastically rational if they have engaged the proper belief formation procedures in arriving at a belief. Say P is best supported by my available evidence. I might come to believe P through hypnosis. Such a belief is propositionally rational, though doxastically irrational.

⁵Notice however, such a presupposition does not entail that a subject’s situation has no bearing on the reasonableness of her beliefs. That Newton did not have access to the kind of information available to contemporary physicists matters in our assessment of him. The bubble chambers, bolometers, and particle accelerators of modern physics give today’s physicists access to data Newton could not have anticipated. That is, his subjective situation, namely what he could not have known, matters in determining whether he was in fact justified in holding to the laws of Newtonian Mechanics.

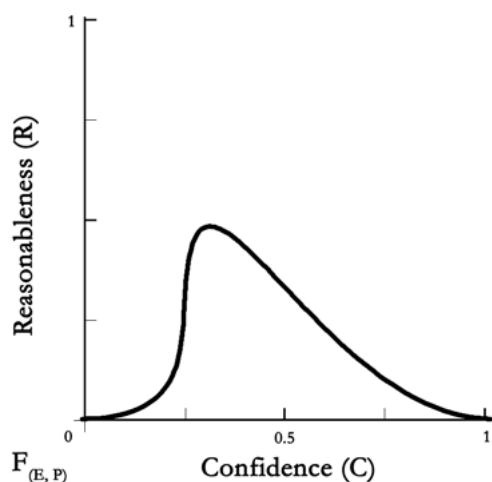
In the disagreement literature in particular, further abstract characterizations of humans *qua* epistemic agents abound. Elga, for example, treats our appraisals of the epistemic abilities of our colleagues as a probability judgment. Specifically, our appraisals are judgments about the probability that our colleagues will be apt in their assessments of evidence in forming true beliefs. Again, this is meant to capture a genuine phenomenon, namely that we trust our fellows to varying degrees. But again, the precision implied by such treatment is not obviously a feature of the phenomenon. The result of this slide into precision is to, in my view, wrongly export puzzles that the formalisms foster to the phenomena.

The proposed solution to the problem of peer disagreement identifies the target of compromise with representations of *reasonableness*, not (first order) weighted beliefs. Consider again our economist concerned with her colleague's registered disagreement. Presented with a body of evidence (E) pertaining to the likely rise in interest rates (P), she arrives at the judgment, a weighted propositional belief ($C(P)$), that interest rates will fall. She believes, at some low level of confidence (C), a certain proposition (P), one expressed by a use of 'that interest rates will rise'. If we hold that there is some non-subjective relationship between an evidence base (E) and the reasonableness of any such weighted propositional belief supported by that (direct) evidence, then the economist is initially justified in maintaining her belief in P , at that weight, if that evidence (most) strongly supports that level of confidence in P . To put matters formally, there is a function $\mathcal{F}(E, C(P)) = R$ that characterizes the actual non-subjective relationship between an evidence base (E), a proposition (P) believed at confidence level (C), and the reasonableness (R) of believing that proposition at that weight. Call the function that best characterizes the reasonableness relation for a given evidence base and range of weighted propositions \mathcal{F}_{ACTUAL} .

The reasonableness relation described suggests a view of epistemic agents as forming beliefs in the process of representing this relation. In forming a belief, we initially consider how a body of evidence pertains to some proposition P , and assess how reasonable one would be in maintaining various levels of confidence in P . Less symbolically, we try to figure out what one could reasonably believe based on some evidence, and how more-or-less reasonable those beliefs are. In making this assessment, we narrow in on the most reasonable position. We can characterize this process in functional terms: we derive what we take to be the characteristic function from the direct evidence to the reasonableness of believing a proposition at various credences, deter-

mine the global maximum value of the function—the confidence level that (by our lights) is most reasonable given the evidence—and settle on this value as our belief. Agents do better the closer their reasonableness representation is to \mathcal{F}_{ACTUAL} . To put matters graphically, our economist presented with some evidence E that pertains to the rise of interest rates (a proposition P) builds a representation indicated by the graph below that plots the reasonableness of believing P (‘That interest rates will rise’) at various levels of confidence, 0–1, concluding with her believing $.3(P)$.

Figure 1: Reasonableness Function



By illustrating in §4 how this shift in target resolves the fundamental problem of disagreement described in §3, I will characterize the process of assessing the reasonableness of a particular belief given a body of evidence as building *reasonableness functions* (*R*-functions for short). As argued above, the very circumstances that give rise to the problem of disagreement presuppose that there is a non-subjective relationship between a body of evidence and what one can reasonably believe based on that evidence. The proposal is that agents are best characterized as building *R*-functions meant to describe this relation. The solution to the *central challenge* of disagreement (described in §3) is that the mere fact of disagreement implores us to revise these representations about the space of reasonable views in accordance with our disputant’s assessment. Sometimes, this requires revising our beliefs, in

cases where our peer’s assessment of the evidence can be gleaned from the indirect evidence. However, as I argue in §4, in most cases of disagreement the mere fact of disagreement fails to supply us with sufficiently rich information about our peer’s reasonableness representation. In these *paradigm* cases one is justified in ignoring their peer’s opinion because they are simply unable to perform the needed compromise. Making that point sufficiently clear requires abstractly characterizing epistemic agents as engaged in the task of somehow deciphering this relationship between reasonableness, evidence, and confidence. For the purpose of clarifying why the weight of indirect evidence varies from context to context, I characterize this task as deriving an *R*-function that assigns reasonableness values to weighted propositions (based on a body of evidence). But such a view is not a psychological claim about the mechanisms of belief formation, and likewise the implied precision of the (loosely sketched) model are not reflective of the muddled and messy way actual epistemic agents approach matters of disagreement. Rather, the purpose is to highlight that in some contexts, but not others, the reasonableness relation is easier to determine, and thereby conveys the pertinence of the indirect evidence provided by a peer.

In the next section I briefly review the Equal Weight view. The purpose is two-fold. First, the EW view serves as a useful case for making explicit both what an account of peer disagreement must accomplish, and what is required in a case of disagreement for accommodation to be feasible. Second, the review highlights how the EW view attempts to solve the *central challenge* of disagreement, and where the view fails. The FEW view aims to correct this failure, by explaining why agents are often justified in ignoring their peers without denying that we have peers in the first place.

3 The Equal Weight view

Our economist is pondering how to address the indirect evidence from her colleague who disagrees with her about future interest rates. One way to resolve her concern is to adopt the Equal Weight view, which yields the normative claim that in cases of peer disagreement we should give our peers’ assessments equal weight to our own, and adjust our confidence in the truth of the disagreed upon proposition to the mean value of our respective weighted beliefs (Elga, 2007). We should “meet in the middle” regarding the veracity of our beliefs (see Elga, 2007; Feldman, 2006; Christensen, 2007). Assuming

that the economist's colleague believes that interest rates will rise with the same veracity as the economist believes they will fall, the EW view demands that the economist should adjust her confidence and become agnostic about the future rise and fall of interest rates.

Our intuitive appraisal of certain cases motivates this view. Consider the following:

Horse Race

Lewis and Max are watching a horse race from nearly identical vantage points. As the race closes, two horses, Quine's Revenge and Holistic Meaning, are "neck-and-neck." As such both Lewis and Max are watching the end of the race attentively. As the horses begin to cross the finish line Lewis comes to believe, based on his observations, that Quine's Revenge edged out Holistic Meaning, while Max comes to believe, based on her own observation, the opposite. Following the conclusion of the race they discuss the results and discover that they disagree.

Were we in Lewis' or Max's position we ought to suspend judgment regarding the outcome of the race. Seemingly, in order for Lewis or Max to maintain their respective levels of confidence about the believed outcome of the race, each would have to hold that they are somehow better disposed to observe the race's outcome. But, assume neither is better positioned.⁶ Intuitively, Lewis should accommodate Max's judgment, and somehow compromise.

Cases like **Horse Race** motivate the *intuition to compromise*, ensconced in the following principle:

- (**Comp**) If an epistemic agent is presented with the results of a peer's deliberation, and that result differs from their own, that agent ought to compromise with that peer.

Max's equally reliable deliberation over the available evidence has yielded a different outcome to Lewis', so we judge that Lewis ought to compromise with Max.

⁶Assume each has perfect vision, that they are sitting in adjacent seats, that they are not intoxicated, and so on. We can assume that neither has a particular interest or bias as to which horse wins—neither has placed a bet, for example. They are equally capable and disinterested in the outcome.

Different views regarding peer disagreement offer an explanation for the plausibility of **(Comp)** by indicating what compromise amounts to, describing the kind of *procedure* an agent must utilize to effect such a compromise and thereby indicating what kind of *information* is needed to follow that procedure. In order to effect compromise in a case of disagreement, one needs to know *what* they are accommodating. Thus, for a compromise to be feasible in a case of disagreement, the indirect evidence must provide the agent with the requisite information.

To see this, consider how the EW view accounts for **(Comp)** in a case like **Horse Race**. The proposed procedure by which compromise is effected is one of averaging: one ought to *average* the *confidence value* of a peer’s weighted belief with one’s own, and adopt that mean value. But to utilize that procedure one must have a particular piece of *information*, the confidence value of their peer’s weighted belief. Lewis naturally has this information, or at least information that reasonably approximates this, because Max provided indirect evidence in their post-race discussion about her weighted belief (e.g. “I completely disagree!”). This indirect evidence proffered the information Lewis needed to perform the averaging procedure, the weight of Max’s belief. We can call such indirect evidence *transparent*:

- (Trans)** An item of indirect evidence is transparent to an agent with respect to a piece of information if that evidence, in an epistemic context, suffices to provide that piece of information to that agent.

If one is expected to incorporate their peer’s assessment of the direct evidence, the indirect evidence espoused by (or about) one’s peer must provide sufficiently rich information about her assessment to characterize the input needed for the compromising procedure. For the EW view, in transparent contexts,⁷ where the features of the disagreement situation combined with

⁷There are various ways in which we might make the notion of transparency more precise. Doing so might make clear that the relation at play is relativized to both contexts and agents. A first approximation might look something like the following:

$$(T) \mathcal{T}_{\langle w, C, a \rangle}(E, I) \leftrightarrow \forall w' [R_C(w, w') \rightarrow (\text{ACCESS}(a, E) \rightarrow \text{KNOWS}(a, I))]$$

This itself is incomplete in ways, since the accessibility relation in R_C would need further articulation depending on how one would capture the notion reflected in the expression “suffices to provide that piece of information.” A subjectivist conception which requires the relevant agent to know the required information in epistemic context C would outline a different set of accessible worlds than a more objectivist account, which would build

the indirect evidence yield information commensurate with a peer's weighted belief, satisfying **(Comp)** simply requires averaging confidences.

However, despite the apparently generality of **(Comp)**, not all cases of disagreement intuitively require that we compromise with a peer. Indeed, *paradigm* cases of disagreement intuitively do not require compromise. Elga (2007) uses the morality of abortion as a paradigm case, where the intuitive response to disagreement is to remain *steadfast*. Finding out that a peer disagrees with me about the moral permissibility of abortion has (it seems) no influence on what I should believe about abortion's permissibility. Faced with disagreement from a colleague they respect, one might react to the lodged disagreement by conferring with them about central arguments, important thought experiments, scientific findings and other aspects of the complex body of relevant direct evidence. But finding that the colleague is equally apprised of the direct evidence, confirming that they are indeed a peer, one is not intuitively compelled to relinquish their pre-disagreement belief. Far from applying **(Comp)**, it seems in paradigm cases of disagreement, in which a complex body of evidence is pertinent to the truth of the disputed claim, we ought to remain steadfast.

Addressing this tension between the compromising and steadfast intuitions is the *central challenge* posed by peer disagreement. The success of any account for peer disagreement hinges on its ability to explain the initial plausibility of **(Comp)**, while curtailing its generality to exclude paradigm cases.⁸ The Equal Weight view accounts for the intuition to compromise by demanding an averaging of credences, often substantially augmenting one's confidence in their position.⁹ More generally, the EW view holds that in

in some normative notion about what a *reasonable* agent would come to know in that context, outlining worlds where the agent (or their counterpart) is ideally rational (say). I leave these technical bits aside, since the intuitive notion is sufficient here make the point about the informativeness of indirect evidence.

⁸One way to accomplish this is to offer a kind of error theory about either the initial plausibility of **(Comp)**, or the intuition to remain steadfast in paradigm cases. Kelly (2010) might be best interpreted as doing the former by insisting that there is no general account of how to incorporate indirect evidence—it is a *mistake* to think that our intuitive response to any case generalizes.

⁹Similar points can be made about cases like the following from Christensen (2011):

Mental Math

After a nice restaurant meal, my friend and I decide to tip 20% and split the check, rounding up to the nearest dollar. As we have done many times, we do the math in our heads. We have long and equally good track records

situations of disagreement we are obligated to give our disputant's opinions a weight that matches the (perceived) probability that they will arrive at the correct answer about the disputed matter, with cases of peer disagreement being those where one takes a disputant to be equally as likely to get the matter correct (Elga, 2007, 490).

But the EW view requires that in *all* cases of genuine peer disagreement that we compromise with our peer by averaging confidences. While this seems intuitive in some cases like **Horse Race**, in paradigm cases of disagreement where we intuitively ought to remain steadfast, the EW view yields the wrong verdict. To the degree such intuitions generalize, the EW view seems to leave these data unexplained. Elga (2007) addresses this initial difficulty, thereby meeting the challenge to constrain the generality of (**Comp**), by distinguishing cases like **Horse Race** from more canonical cases of disagreement, like abortion's permissibility. He insists that in the latter kinds of cases, we have grounds to demote our peer, taking them to be an epistemic *inferior*. Elga's argument is that given the *complexity* of the evidence involved in these paradigm cases, the probability that our peer has (by our lights) incorrectly assessed the evidence is quite high. In this way the EW view explains the *prima facie* generality of (**Comp**), while maintaining that we are justified in remaining steadfast in paradigm cases of disagreement because they are not cases of *peer* disagreement. Paradigm cases do not involve peers at all, and so do not run afoul of a principle like (**Comp**).

However, the EW view suffers from a few counterintuitive implications. First, the explanation offered for remaining justifiably steadfast in paradigm cases plainly denies the basis for the problem posed by peer disagreement. The response insists counter-intuitively that, in paradigm cases, where disagreement matters the most, we have no peers. The very problem of peer dis-

at this (in the cases where we've disagreed, checking with a calculator has shown us right equally frequently); and I have no reason (such as those involving alertness or tiredness or differential consumption of coffee or wine) for suspecting one of us to be especially good, or bad, at the current reasoning task. I come up with \$43; but then my friend announces that she got \$45.

As with **Horse Race**, Christensen and his friend should "split the difference" in their respective beliefs, and remain agnostic about whether each party owes \$43 or \$45. Again, either Christensen or his friend would have to have a good reason for thinking that they have done a better job in analyzing the relevant numbers. But as the case is described both have agreed, from a long established relationship, that they are equally competent, both generally, and in the instance, at performing the required calculation.

agreement is motivated by a laudable humility regarding our epistemic abilities, which the lodged disagreement of our peer serves to remind us about. Elga’s response replaces this humility with hubris. Maybe such hubris is demanded in some cases, but as a general solution it seems unwelcome. Second, the view requires us to be spineless in the face of multiple disputants, and particularly, epistemic inferiors. With enough inferiors that disagree with me, accommodating each in turn mandates that I agree with them, irrespective of their abilities to assess direct evidence.¹⁰ That is, given enough disputants, an epistemic agent would be obligated to agree with them, having her initial assessment *swamped* by their collective judgments. Third, because the EW view gives swamping, out-sized weight to indirect evidence, it permits a kind of bootstrapping from an initially irrational appreciation of the evidence in a given case of disagreement, allowing an agent to rationalize an otherwise irrational position (Kelly, 2010). Kelly (2010) concludes, on the basis of this problem (and others) that we should adopt the Total Evidence (TE) view, which denies that there is a general treatment of peer opinions that holds across evidential contexts. As he puts the point:

On the present view . . . neither [the direct nor indirect evidence] suffices to fix the facts about what it is reasonable to believe . . . Granted that, on the Total Evidence view, both the first-order [direct] evidence and the higher-order [indirect] evidence count for something, which kind of evidence plays a greater role in fixing facts about what

¹⁰See Elga (2007), Christensen (2007), and Sosa (2010) for discussion of this problem. None of these discussions directly address the problem of spinelessness in the face of epistemic *inferiors*. But the problem is easy to see. To spell out the problem: an inferior and I disagree about some body of evidence as it relates to some proposition P , whereby I conclude, with full confidence that P is true— $1(P)$ —and he concludes with full confidence that P is false— $1(\neg P)$ or $0(P)$. Surely my evaluation of my inferior, however daft, ought to give them *some* probability of getting the matter correct. Suppose this probability is 1%, or .01. On the EW view I then should give his assessment regarding P weight in accordance with this assessed probability, and thus should believe, with some small amount of confidence that P is false— $0(P)$ —deflating my confidence in P being true ever so slightly given my correlative likelihood of getting the right answer at 99%. Thus I should adjust the weight of my belief down $\frac{1}{.99/.01}$, or 1/99th to $.9898(P)$. Now, with this new reasonable assessment regarding P based on the evidence, I meet yet another equally inferior evaluator who also believes in the falsity of P , $0(P)$. Again, my confidence in P should slip another bit $\frac{.9898}{.99/.01}$, or .00997 to $.9798(P)$ (rounding a bit). Encountering more inferiors would obligate me to continue to deflate my confidence in P until I agree (more or less) with these inferiors.

it is reasonable to believe?

It is a mistake . . . to think that there is some general answer to this question. In some cases, the first-order [direct] evidence might be extremely substantial compared to the higher-order [indirect] evidence . . . in other cases the former tends to swamp the latter. (Kelly, 2010, 142)

Of course, this response denies the force of the intuition to compromise generated by cases like **Horse Race** (and **Mental Math** in n. 8 above). If there is no general response to indirect evidence that can be abstracted away from the particular epistemic context of a given disagreement, that we feel compelled to compromise in some kinds of cases but not others admits to no general explanation. Such a response should be a position of last resort.

So, if we can offer a view that can explain what the TE view cannot, without falling prey to the counter-intuitive problems that lead to the swamping objection for the EW view, such an explanation should be appealing. In the next section I offer a general strategy that purports to do precisely this. The argument on offer is an appeal to parsimony, offering a view that can account for the intuition to compromise that motivates the EW view, while explaining why we are often unwilling to compromise in cases of disagreement. If we can account for the *prima facie* generality of (**Comp**) without either denying that one has epistemic peers, or succumbing to the swamping problems that plague the EW view, while still giving a general account of how to incorporate indirect evidence in the way the TE view denies, such a view should be seen as preferable to both. In what remains, I motivate an account that accomplishes exactly that by shifting the target of compromise from weighted beliefs to representations of reasonableness.

4 The Functional Equal Weight view

4.1 The FEW View and Uniqueness

The central claim of the FEW view, indicated in §2, is that the primary outputs of deliberations are reasonableness representations. Prior to an encounter with a disputant, epistemic agents somehow represent the reasonableness relation that holds between direct evidence, and weighted beliefs. This process of assessing reasonableness is characterized as producing *R*-functions that represent \mathcal{F}_{ACTUAL} —the characteristic function that describes

the worldly relation that holds between an evidence base and a range of weighted propositions. Construing epistemic agents as if postulating a functional relationship between evidence and reasonableness proffers a variety of resources for capturing our varied intuitive responses to paradigm and compromising cases. We can capture the *prima facie* generality of the intuition to compromise expressed in (**Comp**), much as the EW view does. Agents, when presented with the output of a peer’s deliberation, ought to “meet in the middle” regarding those outputs, giving their peer’s assessment equal weight. One ought to compromise by performing a procedure analogous to aggregating their *R*-function with their peer’s.

The intuition to compromise is grounded in the following thought: given that my peer and I are both equally good at assessing evidence (in the relevant domain), whatever processes we respectively engage to make such assessments are, as best as we can tell, equally reliable. If we construe the output of that process as a weighted propositional belief (as in Elga (2007)), this demand to compromise requires that we aggregate these beliefs by averaging, with the unwelcome result that we cannot remain steadfast because the indirect evidence is given swamping weight. But on the FEW view, the output of deliberation is characterized by a function regarding the reasonableness of various positions in relation to the evidence. The object of compromise, then, is not a (simple) weighted belief, but the assessments characterized by these functions. What peers ought to normalize are these representations that FEW *R*-functions approximate.

To use an example, take the case of **Paradigm** from Kelly (2010).


Paradigm

You and I both accept the Equal Weight View as a matter of theory. Moreover, we scrupulously follow it as a matter of practice. At time t_0 each of us has access to a substantial, fairly complicated body of evidence. On the whole this evidence tells against Hypothesis *H*: given our evidence, the uniquely rational credence for us to have in *H* is 0.7. You give credence 0.7 to *H* while I give it 0.3. At time t_1 , we meet and compare notes. Because we both accept the Equal Weight view, we converge on credence 0.5.

You and I arrive at our respective beliefs about *H*. On the FEW view our beliefs correspond to the respective confidence levels that are assigned the global maximum values of the functions that characterize our reasonableness representations. Suppose then our reasonableness representations

regarding the actual relationship between E and H , are best represented by R -functions with the shape of a normal distribution, but centered around different maximum values. Normalizing these functions would generate the function represented in the figure below:

Figure 2: Aggregating Reasonableness Functions



Normalizing these functions as such, we ought to arrive at the belief of $.5(H)$.

Before showing how the FEW view resolves the central challenge posed by disagreement, a central conceptual issue regarding reasonable beliefs bears mentioning, as our graphs thus far have implied a particular view about what is rationally required by a body of evidence. For a given body of evidence, one might wonder whether there is more than one rational position to hold.¹¹ There are seemingly two views on offer here.¹² First, that there is exactly one uniquely rational position most supported by the body of evidence; a single level of confidence made rational by the available evidence. Second, one could hold that there are (at least sometimes) multiple levels of confidence that are

¹¹I am indebted to an anonymous reviewer at another journal for raising this point.

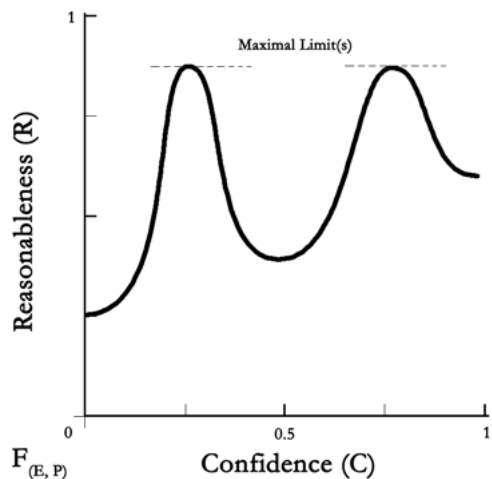
¹²For completeness, one could adopt a third position, that given a body of evidence there are *no* reasonable beliefs to have. Such a view expresses a kind of epistemic nihilism that eschews the problem of peer disagreement from the outset, since no evidence, indirect or otherwise, renders a belief rational.

rational, given a body of evidence. That is, we can ask whether the following is true:

(UT) For a given body of evidence and a given proposition, there is some level of confidence that is uniquely rational to have in that proposition given the evidence.¹³

In terms of reasonableness relations and their characteristic functions, there are three ways of understanding these two positions regarding the *Uniqueness Thesis* in (UT). The first is to say that the global maximum value for every \mathcal{F}_{ACTUAL} is the *only* rational position to hold, given that this is the value for C at which the function is greatest. This view is a fairly demanding. Given a non-subjective relationship between evidence and reasonableness, (UT) is (in part) an empirical claim about disagreement situations; namely that none of the characteristic functions that govern the relationship between a body of evidence and the reasonableness of believing some proposition based on that evidence have more than one maximum value between 0 and 1. That is, the defender of (UT) would have to deny that two C -values could yield the same global maximum R -value of the function. She must deny that any \mathcal{F}_{ACTUAL} could have the following multi-modal character:

Figure 3: Multi-Modal Reasonableness Function

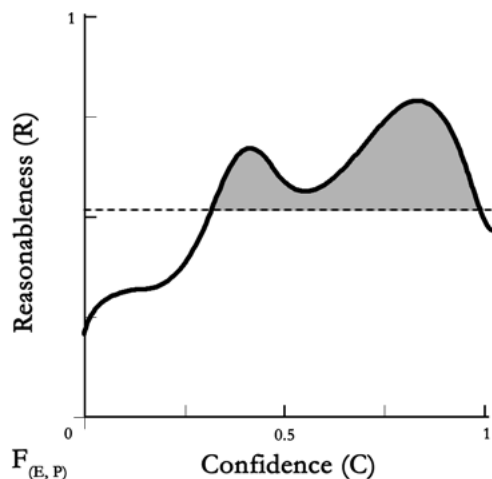


¹³For discussion see Feldman (2007), Christensen (2007), and White (2005).

In the face of situations in which multiple positions seem equally rational, one could maintain the spirit of (UT) by endorsing a second view whereby one is rational in holding a particular (weighted) belief, just in case the confidence level in that belief is *a* global maximum of the characteristic reasonableness function. This leaves open the empirical matter about whether the actual relationship between *some* body of evidence and the reasonableness of believing a proposition with a level of confidence actually has multiple global extrema, which happen to share the same (maximal) value. Again graphically, this in-spirit view would concede that both maximum values in Figure 3 above would be rational. One would be rational in believing $.25(P)$ or $.75(P)$, given that the value of the function at those levels of confidence are both global maximum values. While this approach to (UT) is not quite as demanding as the first staunch view, given that reasonable people on this in-spirit view can reasonably disagree, they seemingly cannot disagree all that often.

The third position on offer is to *deny* the Uniqueness Thesis in spirit, requiring that some other standard be applied to determine which levels of reasonableness render a (weighted) belief rational. We need not settle here what determines that standard, assuming we wish to deny the Uniqueness Thesis. But one obvious solution, given the functional treatment of epistemic agents, is that there is some reasonableness-value that marks the distinction between the rational positions and the irrational positions. If \mathcal{F}_{ACTUAL} for a given context assigns a belief an R -value greater than some demarcating value, then that belief is reasonable. As indicated by the shaded area in Figure 4, an epistemic agent would be justified (or rational) in holding a position regarding P , assuming his knowledge base is relevantly characterized by E , if the characteristic function assigns that belief an R -value that is sufficiently high.

Figure 4: Multiple Rational Values



One virtue of the Functional Equal Weight (FEW) view is that by treating epistemic agents as building reasonableness representations, characterized as deriving a functional relationship of this kind, we are able to quite naturally outline the commitments of various views on theses like **(UT)**. This is helpful both in understanding these views, and in getting clear on what commitments they make. Put another way, we can see that accepting **(UT)** not only belies the notion that there are some matters on which reasonable people can reasonably disagree, but that the supporter of **(UT)** is also committed to a rather strong empirical claim about the distribution of reasonable positions supported by a given body of evidence.

A second (related) virtue of the FEW view is that it straightforwardly accounts for an aspect of our epistemic practices: we routinely make judgments not merely about what to believe, but how reasonable various beliefs are given available evidence. As noted earlier, no one judges Newton to be unreasonable for inventing Newtonian Mechanics as a result of the (now apparent) behavior of quantum particles. Newton's available evidence renders Newton's beliefs (incredibly) rational, despite their falsehood. Likewise, when faced with cases of disagreement with others, invoking a measure of impartiality often presents the positions of others as reasonable. For example, given the current available evidence, it is reasonable to think anthropogenic climate change is not quite as dire as some claim, yet denying that humans

have caused the climate to change is completely unreasonable (Rosenberg et al., 2010).

The central virtue of the FEW view, however, is that it can explain what the EW view cannot, namely why indirect evidence (seems to) deserve more weight in some disagreement situations compared to others. That is, it can resolve the central challenge of peer disagreement in accounting for the intuitive plausibility of (**Comp**), while also curtailing its mandate in paradigm cases of disagreement, all while avoiding the counterintuitive results of the EW view.

4.2 Resolving the Challenge

We saw above how compromise is generally attained on the FEW view in the discussion **Paradigm**. Much like the EW view, the prescription is to find a kind middle point between two assessments. Incorporating a peer's input involves a procedure analogous to aggregating two R -functions. But in accounting for (**Comp**) this way, the FEW view generates the same swamping problems, since aggregating, much like averaging, enables repeated encounters with disagreeing peers (and inferiors) to swamp one's initial assessment of the direct evidence. Aggregating as the FEW view mandates seems to yield the wrong verdict in paradigm cases of disagreement where we intuitively take steadfastness to be justified, in much the way the EW view does.

The FEW view, however, avoids the counter-intuitive implications of the EW view because, as we'll see, in many disagreement contexts one simply lacks the information needed to perform the mandated compromising procedure. In paradigm cases of disagreement where such information is not available, one is justified in remaining steadfast, thereby constraining the generality of (**Comp**) and forestalling the swamping problem. This is accomplished without implausibly denying that one has epistemic peers. In this way, the FEW view is more parsimonious than the EW view, maintaining the virtues and avoiding the vices of the EW view in offering a general account for dealing with peer disagreement. While the FEW view concedes that one ought to compromise by normalizing R -functions, we can explain the intuition to remain steadfast in paradigm cases once we recognize that the indirect evidence in paradigm cases is not *transparent*.

To see this, note that in actual situations of disagreement the indirect evidence from our disputant does not straightforwardly yield information at the level of granularity needed even to assess the precise heft of a disputant's

weighted propositional belief. So, for example, in describing **Paradigm**, to indicate that at t_1 I know you believe $.7(H)$ is to provide a kind of short-hand for the much less clear real-world claim that, I know you believe H more-or-less strongly. Of course, for purposes of abstracting away from real-world difficulties, such stipulations are useful. But, if a theory about the normative force of disagreement is supposed to capture our intuitive judgments about *actual* world disagreements, assuming that paradigm cases of disagreement involve an exchange of information between peers in the form of a precisely weighted propositional belief may assume a bit too much.

If this is right, then considering that, on the FEW view, the correct objects of compromise are reasonableness representations described by R -functions, the level of granularity required in disagreement discoveries to effect compromise is an order of magnitude more detailed than the more modest assumption about weighted beliefs. The information one needs in order to perform the required aggregation must be quite rich, enough so to characterize a disputant's derived R -function. But the indirect evidence from a peer, their stated level of disagreement, does not provide one with that kind of information. Receiving even Bayesian judgment information from a disputant does not amount to receiving information of the kind abstractly characterized by a confidence-reasonableness pair. On the FEW view, one should interpret claims like 'I believe $.8(P)$ ' as the following: 'My assessment of the evidence has led me to characterize \mathcal{F}_{ACTUAL} as having a global maximum at $C = 0.8$.' But knowing that my disputant would assign $.8(P)$ the global maximum value of their R -function does not provide one with knowledge of the particular reasonableness value assigned to $.8(P)$. Thus, in the paradigm disagreement situation involving complex evidence, the mere fact of disagreement does not provide a great deal of information about our disputant's derived reasonableness function. Given that the required normalizing operation should be executed on such functions, in most situations the mere fact of disagreement offers little guidance as to how one should alter their beliefs.¹⁴ The normative force of (**Comp**) is undermined in paradigm

¹⁴Of course, one could incorporate the weighted belief of a peer into one's own reasonableness representation. The formal analog is something like adding that single point to the derived reasonableness function, averaging the disputant's judgment at the relevant point on one's own derived function. That is, one could take the disputant's R -value at his professed most reasonable C -value and assign an R -value for the function at that C -value that is the average of the two R -values. This assumes first that in giving their judgment that my disputant conveys something akin to a Bayesian judgment, and second

cases by the fact that one simply cannot perform the mandated aggregation, since paradigm cases involve indirect evidence that is non-transparent with respect to a peer's reasonableness representations.

And this is true regardless of what aggregation operation turns out to be best. The aggregation strategy implied by the graph in figure 2 above is a point-by-point averaging. If epistemic agents always derived normally distributed functions (with similar standard deviations) such an averaging method for accommodating the intuition to compromise would be uniformly applicable to all disagreement situations. But to assume that epistemic agents would always derive normally distributed functions is unreasonably optimistic. For one, the \mathcal{F}_{ACTUAL} functions of many disagreement situations might well be non-gaussian. There is no *a priori* reason to think the actual relationship between confidence levels and reasonableness (given an evidence base) will admit to a normally distributed treatment in all cases of disagreement.¹⁵ Second, even if the actual relationships in the world were normal, that alone does not license the assumption that epistemic agents will build reasonableness representation that reflect such normal distributions. That is, granting that real-world distributions *are* normal is not itself a reason to assume that epistemic agents will *think* such distributions are normal. In assessing the confidence-reasonableness relationship, any given agent might derive any sort of function: one with a flat distribution, a disjointed function, a partial function, etc. Given the diversity of possibilities, an agent is left in a particularly impoverished position about their peer's R -function. The indirect evidence in such cases is *not transparent* with respect to the information needed for aggregating.

This leaves open a rather important aspect of the FEW view, namely, the kind of aggregation operation an epistemic agent ought to perform to accommodate the intuitions to compromise. Investigation into the nature

that his judgment also conveys the R -value associated with that Bayesian judgment. But even if we accept these assumptions, and perform the averaging operation, the resulting function would likely have the same extrema, at least in cases where the initially derived R -function has a maximum value that is somewhat high, or fairly high relative to the rest of the function.

¹⁵These actual relationships *might* often be characterized by normal distributions, given that normal distributions (are thought to) naturally exist in abundance; cf. the Central Limit Theorem. But not only is there no reason to think all such relationships will abide by this trend, there may well be reason to think that the normal distribution is not as common (or normal) as we think; see Lyon (2014) regarding the frequency and normalcy of normal distributions.

of the correct aggregation operation is indeed needed to fully flesh out the FEW view, and will likely be a difficult question to navigate in ways that go beyond the fundamental problems of judgment aggregation (List & Pettit, 2002). For one, empirical investigation into our capacities in this regard is required. There may well be cognitive limits on the kinds of operations humans can perform in aggregating over the indirect evidence. Thus even if there were some way to determine the most appropriate means of aggregating R -functions, there may well be limits on the ways the human mind can manipulate the reasonableness representations these R -functions characterize which precludes certain strategies and favors others. But this paper will not attempt to offer an answer to this difficult question. Importantly, this concern is orthogonal to the following main point: regardless of which aggregation operation is the most appropriate, *any* such strategy would require that the functional inputs to this operation are robustly characterized. That is, to perform *any* operation over my reasonableness representation and that of my disputant I'd need to know (to some large extent) the shape of my disputant's representation.

Most disagreement situations cannot be fairly characterized as an exchange of information commensurate with the detailed kind required to perform the analog operation of normalizing two R -functions—the operation needed to satisfy our compromising intuition in **(Comp)**. What I must get from my disputant, and what I rarely in fact get, is information about my disputant's representation of reasonableness at the granularity of a derived R -function. The generality of **(Comp)** is constrained as a result, since agents in paradigm cases of disagreement cannot perform the mandated aggregation, because they lack information that determines their peer's R -function. Nor do they receive sufficiently rich information that approximates this, at least not in the typical case of disagreement.

But not all cases of disagreement are like **Paradigm**, where the intuitive response is to remain steadfast. As we saw in §3, cases like **Horse Race** intuitively require that we compromise with our peers, not ignore their registered disagreement. Importantly, however, in these contexts the indirect evidence *is* sufficiently transparent to perform the mandated aggregation. The cases that motivate the intuition to compromise indicate contexts in which we have pre-evidential reasons for thinking that \mathcal{F}_{ACTUAL} has a particular shape. For example, in **Mental Math** where two friends disagree about the amount that each owes on a restaurant bill—even after repeated calculation—we are compelled to think each should relinquish ground in the dispute (from

Christensen, 2011, see n. 8 above). But, the reason mathematical-logical reasoning is so attractive is because problems can be given formulaic solutions. The rules of mathematical systems, once we get clear about the kind of operations required for the calculation, *determine* the space of acceptable solutions. In cases like **Mental Math**, they determine a unique solution. As such, in assessing what beliefs are reasonably supported by the evidence, there is (essentially) only one reasonable position given the evidence, with all others seemingly unjustified by the evidence.

Given the character of mathematical-logical contexts of disagreement, we have general knowledge about the typical relationship between mathematical evidence and reasonableness—knowledge we bring to bear on the informativeness of our peer’s registered disagreement. In such cases, acquiring information that approximates the Bayesian judgment of my peer essentially describes their representation of what beliefs are reasonable—put in FEW terms, the analogous R -function is quite boring. Given the kind of knowledge one brings to mathematical disagreements prior to any specific information about the case, when this knowledge about the shape of \mathcal{F}_{ACTUAL} is combined with the judgment of our peer, the result is the kind of information lacking in most real-world cases of disagreement like **Paradigm**. In effect, when Christensen and his friend confer after they perform their respective calculations in **Mental Math**, they provide each other with their weighted propositional belief, which, when combined with the pre-evidential information we bring to mathematical cases, effectively provides a characterization of their respectively derived R -functions. Thus, in cases of this type we *can* perform the normalizing operation that our intuition to compromise demands, because the indirect evidence is transparent. For these reasons such cases initially seem to support the EW view, since learning about a peer’s weighted belief effectively conveys the contours of their reasonableness representations.

Perceptual cases like **Horse Race** have a similar character. The relationship between the available evidence and reasonableness is fairly determinate. Given the biological facts about vision and intraspecies variation, the likelihood that two different people given identical stimuli will have contentfully distinct representations is fairly low. While most epistemic agents do not bring a substantial body of technical information about human perceptual systems with them to perceptual cases, they do seem to bring with them less technical first-hand knowledge that approximates this. Of course differences in stimuli can produce differing representations. But if two perceivers are situated in similar positions, paying attention to the perceived scene at simi-

lar levels, they are likely to have highly similar phenomenology. In the cases that motivate the intuition to compromise in **(Comp)**, these parameters are stipulatively fixed. Thus, holding fixed the stimuli, the representation of two observers should be quite similar. To motivate the view that two viewers of a scene with nearly identical phenomenology of the scene could reasonably come to different conclusions about what happened, one would have to imagine that there is some complexity in interpreting this phenomenology—a complexity that one could be better or worse at assessing. On some occasions, there is such complexity. Whether or not a particular object counts as a ‘chair’ might involve a kind of complex interpretation of some visual representation, commensurate with the complexity (or vagueness) of notions like CHAIR. But importantly, in the kinds of visual cases that motivate the compromising intuition, this complexity is absent. What distinguishes a HORSE-RACE-WINNER is clear in a way CHAIR is not. As such, the stimuli in such cases determine what one can reasonably believe. Much like the mathematical case, we have pre-evidential information about the shape of the reasonableness representation: one position will be very reasonable, and the others will be quite unreasonable. Thus, learning about my peer’s weighted belief in perceptual cases is sufficiently informative about the character of their derived *R*-function such that one can perform the required compromising operation. To put the point less technically, perceptual cases are of the sort that antecedently we think have clear answers, in much the way disputes about (say) abortion do not.

4.3 Justifying the Steadfast Response

To summarize, the FEW view achieves what the EW view cannot. The *prima facie* generality of the intuition to compromise in **(Comp)** is accounted for without carrying the counter-intuitive baggage of denying that one has peers. A small measure of epistemic humility suggests that at least some of those with whom we disagree are equally suited to assess the available evidence. The EW view adopts this unwelcome commitment in order to address the intuition to remain steadfast in paradigm cases, with the hope of avoiding the swamping problem.

The FEW view is faced with a similar concern, insofar as abiding by **(Comp)** *prima facie* requires aggregating one’s reasonableness assessment with their peer’s. As such, meeting more peers who provide me with their assessments, characterized by the form of *R*-functions, requires aggregating

each of these functions in turn, swamping my initial assessment. But given that paradigm cases of disagreement fail to provide me with the information needed to perform the mandated aggregation, I am justified in ignoring, but not downgrading, my peer and remain steadfast. In disagreements where the indirect evidence is transparent (regarding reasonableness representations) aggregation is both possible and mandated. But because the indirect evidence in paradigm cases is *non-transparent*, aggregation is not possible. The normative weight of the intuition to compromise still applies because one has indirect evidence from a genuine peer, but in paradigm cases agents are simply unable to perform the procedure compromise demands.

However, it does not simply follow from an agent's inability to perform the needed aggregation that they are thereby justified in doing *nothing* in response to a peer's registered disagreement. After all, if there is some procedure that an agent is able to perform which yields the same, or sufficiently similar result, then they ought to enact that procedure. If I see a child drowning in a pond that I know I cannot reach in time, this alone does not justify my complete inaction in letting the child drown. If an inattentive bystander is in a position to pull the child from the lake is within shouting distance, my remaining silent is despicable. Only if there is *no* action available to me that I would reasonably expect to result in the child's survival am I justified in ignoring the child's plight.

To justify the steadfast response on the FEW view, there must be no plausible alternative to the mandated aggregating procedure available. That is, we must show that there is no procedure that an agent can enact in a paradigm case of disagreement that will produce the same, or similar results to the aggregating procedure mandated in transparent contexts. A completely decisive argument for this claim would require showing that *every* possible alternative procedure available to an agent in paradigm cases of disagreement fails in the relevant respect. Lamentably, proving negatives is notoriously elusive. Instead, I offer an argument from analogy, illustrating that the task of finding an alternative procedure in paradigm contexts is unlikely to succeed by showing how difficult a similar, and simpler task is.

Respecting the intuition to compromise on the FEW view involves a procedure commensurate with aggregating *R*-functions, assuming the indirect evidence is transparent with respect to a peer's reasonableness representation. In paradigm contexts, the lack of such transparency serves as a barrier to compromise. To see how profound this barrier is, and why it obstructs any path to an alternative procedure, let's consider a world that is non-

transparent with respect to the weighted *beliefs* of others.

Suppose we were in a particularly impoverish epistemic world, at least with respect peers' assessments regarding the matters that concern us. Whenever a peer lodges their disagreement at this world, thereby providing indirect evidence, this evidence (in that context) does *not* convey information that approximates their weighted belief. In this world it is as if agents are only capable of indicating to others *that* they disagree, without being able to convey *how much* they disagree.

Likewise, let's assume for the sake of the analogy that the EW is correct—explaining the intuition to compromise in (**Comp**) is had by averaging the weighted beliefs of a peer with our own. Taking these two suppositions together, the information one needs for compromising at this world is a peer's weighted belief that no indirect evidence ever provides. Because this is a non-transparent world regarding weighted beliefs, no amount of indirect evidence that a peer provides contains the information needed for averaging credences in the way compromise demands in transparent contexts.

Let's consider then a case like **Horse Race** at this impoverished world. Lewis comes to believe strongly that Holistic Meaning won the race, believing $.8(P)$. He then learns that Max, a peer, disagrees. Given that the EW view is (by hypothesis) correct, Lewis ought to average credences with Max, meeting in the middle. But given that this indirect evidence is non-transparent, all that Lewis knows is *that* Max disagrees with him. So he does not have the information he needs to fulfill the EW demand to average credences, as he has no second credence with which to average. How then should Lewis accommodate the indirect evidence? What *alternative procedure* ought Lewis utilize?

A first suggestion might be that Lewis can treat Max's belief as *polarized*, being either $1(P)$ or $0(P)$. The alternative procedure is still a kind of averaging operation, but utilizes information that is (presumably) available to Lewis given his opaque indirect evidence. Max says she disagrees, so Lewis takes that to indicate that Max disagrees with him as much as possible. There are two immediate problems with this response. First, the polarizing strategy assumes Lewis has more information than he has. All Lewis knows in the impoverished world is that Max disagrees with him. Considering that the results of deliberation are taken to be weighted beliefs, disagreement itself is nothing more than a difference in credence between two agents' beliefs. By stipulation, *all* that Lewis knows is that Max has a different weighted belief, which is consistent with *any* confidence value that isn't 0.8. For all that

Lewis knows given the indirect evidence, Max believes $.9(P)$. Treating Max's registered disagreement as if she believes $0(P)$ upon learning the mere fact that Max disagrees, would reach a dramatically incorrect outcome. Lewis' polarized averaging would result in him believe not- P (i.e., $.4(P)$), whereas were the indirect evidence transparent, Lewis ought to be *more* confident in P (again according to the EW view).

But let's suppose that Lewis' position is not *that* impoverished with respect to Max's weighted belief. So let's assume that the world is less impoverished, and Max's registered disagreement at least tells Lewis that Max believes not- P . When Max indicates that she disagrees with Lewis, this indicates that her credence is some value less than 0.5. As such, polarization at least tracks whether Max would assent to P (or not). This leads to the second problem with the polarization procedure. Polarization, even with our more transparent evidence, would mandate that nearly all disagreements with peers result in excessive accommodation. Of course this is most flagrant in cases where peers do not disagree that much. If Max came to believe that Holistic Meaning lost, but was not confident in that conclusion, believing $.45(P)$, Lewis' position, post the polarizing procedure, offers a very different result than the transparent rational outcome—which would have Lewis believing $.625(P)$. He would be obligated to abandon his belief that Holistic Meaning won, contrary to what the transparent outcome would license. And of course, Max would also over-accommodate, with the odd result that peers nearly always swap beliefs. Applying the polarizing procedure across multiple contexts would lead not only to situations in which Lewis adopts the incorrect credence, but in which he routinely adopts the wrong belief (e.g. believing not- P when believing P is transparently mandated).

These considerations suggest a second option for an alternative to the averaging procedure mandated in transparent contexts, namely one that makes reasonable assumptions about the *distribution* of confidence levels that peers arrive at across cases of disagreement. On this proposal Lewis is justified in making an educated guess about Max's confidence in P based on more general disagreement trends, and then averaging based on those assumptions. Lewis is informed that Max, a peer, disagrees. He might suppose that this evidence would only be made available if Max *genuinely* disagreed. Insofar as Max would reject P , Lewis is justified in taking Max to have low confidence in P , (say) at least as low as $.25(P)$. Further, Lewis reasons that in cases of genuine disagreement, confidence levels are normally distributed around a substantively low confidence. Most cases of peer disagreement involve a

confident, but not overly confident, assessment of the direct evidence, and so disputants tend to arrive at significant, though modest levels of confidence. So, the argument goes, Lewis should adopt the following strategy:

(Mean) Across situations of peer disagreement, take a peer's registered disagreement as indicating either a belief of either $.75(P)$ or $.25(P)$.

Depending on whether one accepts P or not, they should interpret the indirect evidence as informing them that their peer's weighted belief is $.25(P)$ if their peer rejects P , or is $.75(P)$ if the peer accepts P .

The alternative procedure in **(Mean)** does a bit better than the polarization procedure. Polarization fails for the obvious reason that the provided evidence of a disagreement is consistent with innumerable confidence values, even granting the assumption that the confidence levels accord with the polarity of the peer's assessment. Over-accommodation ensues. The proposal in **(Mean)** will also require over-accommodation, but reduces the degree of excess accommodation to mirror the degree that the case deviates from the (assumed) mean. For all Lewis knows he's in a disagreement that is an outlier, deviant from the mean. In such cases Lewis will sometimes adopt a belief contrary, and thus insufficiently similar, to the transparently mandated one. But, as a defender of **(Mean)** will point out, longitudinally this strategy will get closer to the transparent outcome, the one that averaging would reach if the indirect evidence was transparent—at least more often than not. Given a normal distribution of confidence levels across disagreement cases, the outlier cases will be rarer than more standard values. So while for any given disagreement the **(Mean)** result might be insufficiently similar to the transparent outcome, overall outcomes will be similar to the ideal transparent ones.

However, the longitudinal point reveals the fundamental problem facing **(Mean)**, and seemingly any alternative procedure: if indirect evidence is not transparent agents are simply not justified in making claims about the distribution of assessments across disagreements. In the case of weight beliefs, the plausibility of assuming that a registered disagreement reflects a modest level of confidence is predicated on having past data that reflects such a distribution. While in the actual world this assumption may seem more or less innocent, in Lewis' world he has no such information. After all, Lewis' position is impoverished with respect to *any* peer's weighted beliefs because indirect evidence is not transparent in *any* context. Lewis has *never*

received information about the weighted beliefs of peers, even when they register their disagreement. Any claim about the distribution of confidence levels across disagreements is based on a sample size of zero, rendering *any* guess about that distribution unjustified. Put another way, Lewis can only make *uneducated* guesses about Max’s weighted belief.

But the problem with (**Mean**) generalizes. Lewis’ task is to take the coarse-grained information offered by the indirect evidence, and somehow reliably convert that into something more precise that approximates the fine-grained information needed to approach the ideal transparent outcome. He must somehow leverage the other information available to him in the disagreement context to accomplish this task. But any attempt to do this would require data regarding the correlation of these contextual factors and a peer’s weighted belief. For example, for Lewis to leverage pragmatic features of Max’s utterance, like that she said “I highly disagree” and not merely “I disagree”, Lewis would have to know that this difference corresponds with a difference in credence. Of course in the actual world this difference seems obvious, but in our impoverished possible world Lewis lacks past data about such correlations. Given that in our stipulated non-transparent world no one is ever aware of the weighted belief behind a peer’s register disagreement, the prospects of offering an alternative procedure to the confidence averaging procedure licensed in transparent contexts seem dim.

Clearly, the actual world has us in a better situation than Lewis’ with respect to the weighted beliefs of our peers. Expressions of disagreement seem informative enough, indicating at least an approximate level of confidence. That is, if the information needed to respect the intuition to compromise is the weighted beliefs of our peers, the actual world is sufficiently transparent. However, indirect evidence in the actual world is regularly *not* transparent regarding the information required to aggregate *R*-functions, the target of compromise on the FEW view. We are, in paradigm cases of actual disagreement, in a position analogous to Lewis’, not with respect to weighted beliefs, but with respect to reasonableness representations. We know what information we need to compromise, but are in an epistemically impoverished situation regarding that information. As such, much like Lewis, we need some alternative procedure that makes use of attainable information in lieu of the information required to engage the aggregating procedure licensed in transparent contexts.

Lewis’ prospects for finding an alternative seem quite dim. The prospects for our actual analogous task are dimmer still, for two reasons. First, our

challenge is an order of magnitude more demanding. Lewis, given his evidence, needs information that merely approximates a value along a single dimension, a confidence interval. The information needed to yield a result close to the transparent outcome in our actual task is comparatively like a binomial function, adding an additional degree of complexity. The problem is an order of magnitude more difficult. Second, unlike Lewis, we do not know the precise nature of the operation that needs to be performed, even in a less impoverished world. On the FEW view, respecting the intuition to compromise in (**Comp**) involves a procedure that approximates aggregating R -functions. But given that there are various competing methods for how that is best accomplished, our task of finding an alternative to the unspecified operation that effects a sufficiently similar outcome is far less clear. Lewis at least knows that he needs to average confidence levels. Given both the comparative complexity of our actual task, and the lack of an explicit procedure to approximate, if Lewis' task seems unlikely to succeed, our actual task should seem even less promising. Of course, this does not show that there is *no* possible procedure that meets the similarity mandate, but it suggests how unlikely we are to find one. If no alternative is available, then one is justified in ignoring a peer's registered disagreement.

Before turning to some general objections, it's worth highlighting one important way in which Lewis' impoverished situation is different from ours, according to the FEW view. In Lewis' world he *never* encounters information that approximates the kind he needs. He never learns the heft of a peer's weighted belief. In the actual world, according to the FEW view, we do *sometimes* get sufficiently rich information from the indirect evidence to effect a compromise. After all, in cases like **Horse Race** we can satisfy the intuition to compromise because the indirect evidence, in concert with a context that constrains the space of possible reasonableness relations, provides information that closely approximates the characteristic R -function of our peer's reasonableness representation. Actual agents thereby have *some* information about past disagreement situations in a way that Lewis does not. One might then worry that the argument from analogy is thereby troubled.

However, this information proves useless in aiding our task of finding an alternative procedure that can be adopted across disagreement contexts. Cases like **Horse Race** are unique contexts, for the reasons indicated above. They constrain the space of reasonableness relations because the evidence in those cases is simple in a manner not true of paradigm cases of disagreement. This simplicity is the reason learning the heft of a peer's belief is sufficient to

justify claims about their reasonableness assessment—they take that belief to be the only sufficiently reasonable option. But this simplicity is exactly what is lacking in paradigm contexts, rendering the information provided, our peer’s weighted belief, insufficiently rich to glean the character of their reasonableness representation. As such, any data one could utilize from past disagreements regarding characteristic *R*-functions would not generalize to the paradigm case. Absent such data, in paradigm contexts we are left in an impoverished position regarding our peer’s assessment, and are justified in ignoring it.

Considerations of parsimony then favor the FEW view insofar as it can explain central intuitions across various cases without taking on implausible commitments. In paradigm cases of disagreement (e.g. the permissibility of abortion) a steadfast response to a peer seems justified. The explanation offered by the FEW of such cases is that, given the complexity of the evidence involved, one simply has no clear means of incorporating the indirect evidence. In contrast, in cases like **Horse Race** we have a clear idea how to incorporate the indirect evidence, explaining why we intuitively take such cases as demanding compromise. Treating the object of compromise as a reasonableness representation makes sense of these responses by highlighting differences in the informativeness of our peer’s registered disagreement across contexts.

5 Objections

Objection 1

In section §2, views like the EW view were criticized because they imply an unrealistic precision to the confidence levels that actual agents assign to beliefs. Credences are more precise than psychological levels of confidence. But now we’re told by the FEW view that agents not only assign precise credences to beliefs, but likewise assign a reasonableness value to those credence-proposition pairs in the form of an *R*-function. Surely then, if the EW view is troubled because of an excess in precision, and the computational demands that it implies, the FEW view suffers worse on this score.¹⁶

¹⁶My thanks to an anonymous reviewer at another journal for raising this concern.

This criticism might be apt if the FEW view proposed that epistemic agents actually calculate reasonableness functions of the sort indicated by the figures above. But the FEW view is far less imposing. The argument makes two assumptions: first, that there is some non-subjective relationship between the reasonableness of believing a proposition at a particular level of confidence, and a body of evidence; and second, that epistemic agents *somehow* represent this relationship, at least in part. Given those assumptions, the response to the problems raised by peer disagreement is simply to adjust the target of the normative force contained in the intuition to compromise. What epistemic agents ought to adjust in light of a disagreement with a peer is their representation of this reasonableness relation. Putting this point in terms of *R*-functions is not an exercise in armchair psychology. Rather, it is merely a means of indicating why the mere fact of disagreement provides sufficient information to guide agents in accommodating a peer in those cases that motivate the compromising intuition, but not others.

Objection 2

About those graphs. The functions visualized therein seem to share very few properties, and thereby indicate very little about what constitutes the real-world relation the FEW view assumes. Indeed, very little is said about what constrains these functions. In the absence of a more explicit delineation of the purported model it's not clear what predictions the view makes, if any at all.

This objection rightly notes that the sketch above fails to say much about the properties of *R*-functions, or the character of reasonableness relations across contexts. There is a good reason for this. The guiding assumption for the FEW view is that there is a non-subjective reasonableness relation that agents represent. We have good reason to think we have such representations. Take, for example, what is required for success on a standard false belief task (Wimmer & Perner, 1983). Representing what an agent (likely) believes when the evidence available to them is no longer accurate requires assessing what a reasonable (or normal) epistemic agent will do with a body of evidence, not merely how they will behave given how the world is. In the standard false belief task, this requires predicting how someone will behave given their misleading information—for example, where they will look for a desired object that has been surreptitiously relocated in their absence. Unsurprisingly, adult performance is remarkably stable (Wellman et al., 2001).

But even children under 4 years old are able to recognize that ambiguous information can be misleading (Lewis et al., 2012). In both cases, the behavioral predictions that adults and children make require that we somehow internalize the relationship between available evidence and the reasonableness of various beliefs.

However, the source and structure of these representations is unclear. After all, how we assess evidence is sensitive to various kinds of bias, in complex and diverse ways (Boyer, 2000; Pohl, 2017). We should expect that our representations of that relationship will be likewise complex. So while we have good reason to posit that we represent the non-subjective reasonableness relation, how this manifests in epistemic agents is far less clear. The argument above remains agnostic on this score as a reflection of this uncertainty. Rather, the view only assumes that, whatever those representations are like, at least one confidence level wins out. This fact is captured by the maxima of an R -function. But the view does make (at least modest) predictions. In whatever way actual agents manage to represent the reasonableness relation, in disagreement contexts where the shape of that relation is clear, we will have strong intuitions to compromise.

An analogy may be helpful. Humans seem able to determine the trajectory of inanimate objects moving through the air. If a billiard ball is dropped from a few meters off the ground, a sighted human paying attention nearby will have little trouble predicting where it will land. If asked to catch the ball mid-flight they will likely succeed. Somehow, they represent the ball's likely trajectory, and those representations can be used in executing motor plans to intercept the ball. But that same capacity to represent the relationship between the external environment, an object, and one's own location in space will likely fail if more environmental variables are introduced. If we introduce an object that responds more variably to the environment, as say a leaf might, we would predict poorer performance in intercepting the leaf compared to the ball. So while there is an objective relationship between the environment and the object, our capacity to represent this relationship in sufficient detail is outstripped by that complexity in contexts like those involving excess wind and leaves. Thus we have a prediction: in environments with fewer variables, where this falling relation is less complex, we would predict greater success in intercepting the object in free fall.

Similarly, in epistemic environments where the relationship between evidence and belief are straightforward, our representations of the reasonableness relation will be more robust and reliable. Given only a few numbers,

and a percentage to calculate, the range of reasonably calculated results is quite narrow, in much the way that the predicted possible locations of the falling billiard ball are limited. So when we have evidence that our predictions are mistaken in such contexts the normative impulse is to adjust the representation on which our judgments rely. But just as we do not need to know how humans represent the predicted movement of free falling objects in order to accurately predict that they will do better with falling metal spheres than cherry blossom leaves, we do not need to know how reasonableness relations are represented to accurately predict which disagreement contexts will manifest intuitions to compromise.

To push the analogy, suppose that when asked to catch the falling billiard ball our onlooker was wearing wedge prism glasses (colloquially “drunk goggles”), whose lenses obscure the visual field by displacing visual objects. On their first attempt to catch the ball while donning this eyewear, they will no doubt fail. But on repeated attempts, they will quickly accommodate the displacement.¹⁷ Given the evidence of their inaccuracy (i.e. failing to catch the ball), they adjust their representation of the possible location of the billiard ball mid-flight. Notice, that an analogous situation with respect to a falling leaf on a windy day will not end with any accommodation. The complexity of the situation outstrips their ability to accommodate the displacement because they cannot isolate whether a given failed catch is because of their displaced vision, or their (muddled) representation of the possible positions of the leaf.

We can think of this evidence of error in the falling projectile case as akin to the lodged disagreement in my peer. In **Mental Math**, where the relationship between evidence and belief is fairly simple, that my peer has disagreed with me compels me to rethink which calculated results are reasonable. So I should adjust my reasonableness representation to accommodate. But in contexts where the reasonableness relation is more complex (or where more variables relevant to the relation are in play as a result of the context’s complexity), accommodating my peer’s lodged disagreement is not possible for reasons analogous to predicting the falling leaf’s location under displaced vision.

¹⁷In experiments where participants threw objects at a target with wedge prism glasses equipped, participants were able to accommodate the visual displacement imparted by the lenses within 10–30 throws, and return to their pre-eyewear accuracy (Martin et al., 1996).

Objection 3

The argument to ignore our peer's assessment in paradigm contexts suggests that when there is no alternative procedure available to respect the intuition to compromise we are justified in ignoring the indirect evidence. But *ignoring the evidence* is a kind of procedure. As such, the argument in favor of that procedure over others would require showing that the ignoring procedure is better than the others available to agents in paradigm contexts. And the argument above does not show that.

Such a reply mischaracterizes the process of assessing evidence, insofar as it fails to distinguish between *how* we (ought to) process evidence, and *which* evidence we should process. Suppose in a timed game of chess Yuki is given overwhelming evidence at the beginning of the match about his opponent's future position, that on the 37th move of the match her queen's knight will land on a particular square, say h6.¹⁸ How should Yuki make use of this information to guide his opening moves? Yuki could attempt to calculate all the various lines that are consistent with that move, and pick the line that is most favorable to him. But even if Yuki had both the super-human ability to keep in mind these numerous possibilities, and the statistical information about which is the most favorable among that enormous set, all of that would take a great deal of time. Assuming, like most of us, that Yuki cannot entertain by brute calculation these various lines and probabilities, the information provided by this evidence is useless. Of course, it is not useless because it is irrelevant—the primary evidence that matters in a chess match pertains to an opponent's likely future moves. The information is useless because he has no way of implementing a procedure that incorporates this information in building a strategy. So Yuki rightly ignores this indirect evidence. In doing so Yuki is not thereby adopting a novel procedure for dealing (or not dealing) with this piece of evidence. He is utilizing the procedure he typically does in developing and altering his strategy, one that does incorporate evidence about his opponent's (likely) future moves. Rather, Yuki is simply failing to take the evidence into account, despite its clear relevance.

On the FEW view, we are often in Yuki's position with respect to the

¹⁸This need not involve any claims about determinism, or paradoxes of free will. Suppose this (probably bad) chess player, for whatever reason, always orients her pieces to ensure that the 37th move has this result.

indirect evidence. There is a procedure available for dealing with indirect evidence, but in paradigm contexts that evidence is useless. Our optimal procedure, yielding the correct result in transparent contexts, cannot make use of the evidence provided. However, ignoring that evidence is not a procedure for incorporating that evidence, any more than Yuki's ignoring the evidence of his opponent's future move requires that he adopt of new procedure for strategy building.

6 Conclusion

The intuitive push to reject both the Direct Evidence view and the Indirect Evidence view likewise motivates us to both accept the EW view and reject it. We seem committed to the primacy of direct evidence in justifying beliefs. However, the fact that someone we revere for their epistemic credentials disagrees with us seems to carry some evidential weight. Certain cases exploit this second intuition, and defenders of the EW view point to these cases as evidence for their view. But the EW view demands that we alter our beliefs in intuitively unreasonable ways, as highlighted by the swamping objection. The FEW view accounts for these competing intuitions. If we view epistemic agents as attempting to codify the non-subjective relationship between a body of evidence and the reasonableness of various beliefs considered in light of that evidence, our intuition to compromise with our peer can be understood as a demand to normalize our respective reasonableness representations. In typical cases of disagreement, however, we are not sufficiently informed about our peer's assessment of the non-subjective relationship captured by her R -function. Merely learning that my peer disagrees with me, even if we grant that by registering her disagreement my peer conveys information at the level of specificity akin to a Bayesian-style belief attribution, does not sufficiently narrow the space of possible R -functions that characterize her assessment of the relevant non-subjective reasonableness relation. As such, I am not sufficiently informed about my peer's assessment to perform the required compromise. While there are some cases in which we feel compelled to compromise, in these cases we bring pre-evidential knowledge which constrains the possible assessments of the reasonableness relation, thereby rendering the indirect evidence transparent. This knowledge explains why learning that my peer disagrees with me is sufficiently informative to perform the required compromise, and why we feel intuitively compelled to compro-

mise in such cases. For these reasons, the Functional Equal Weight view both avoids the problems of the Equal Weight view and accommodates the intuitions that initially motivates the view. Additionally, the Functional Equal Weight view explains what the Total Evidence view cannot, namely why we give indirect evidence differing weight across (types of) disagreements.

References

- Boyer, P. 2000. Natural epistemology or evolved metaphysics? Developmental evidence for early-developed, intuitive, category-specific, incomplete, and stubborn metaphysical presumptions. *Philosophical Psychology*, 13(3), 277–297.
- Christensen, D. 2007. Epistemology of Disagreement: the Good News. *Philosophical Review*, 116:2, 187–217.
- Christensen, D. 2011. Disagreement, Question-Begging and Epistemic Self-Criticism. *Philosopher's Imprint*, 11:6, 1–22.
- Elga, A. 2007. Reflection and Disagreement. *Nouûs*, 41, 478–502.
- Elga, A. 2010. How to Disagree about How to Disagree. In R. Feldman, & T. A. Warfield (Eds.) *Disagreement*, (pp. 173–186). Oxford University Press.
- Feldman, R. 2006. Epistemological Puzzles about Disagreement. In S. Hetherington (Ed.) *Epistemology Futures*, (pp. 216–236). Oxford University Press.
- Feldman, R. 2007. Reasonable Religious Disagreements. In L. Antony (Ed.) *Philosophers without Gods: Meditations on Atheism and the Secular Life*, (pp. 194–214). Oxford University Press.
- Feldman, R., & Warfield, T. 2010. Introduction. In *Disagreement*. Oxford University Press.
- Kelly, T. 2010. Peer Disagreement and Higher-Order Evidence. In R. Feldman, & T. A. Warfield (Eds.) *Disagreement*, (pp. 111–174). Oxford University Press.

- Lewis, S., Lidz, J., & Hacquard, V. 2012. The Semantics and Pragmatics of Belief Reports in Preschoolers. In *Semantics and Linguistic Theory*, vol. 22, (pp. 247–267).
- List, C., & Pettit, P. 2002. Aggregating Sets of Judgments: an Impossibility Result. *Economics and Philosophy*, 18, 89–110.
- Lyon, A. 2014. Why are normal distributions normal? *The British Journal for the Philosophy of Science*, 65(3), 621–649.
- Martin, T., Keating, J., Goodkin, H., Bastian, A. J., & Thach, W. 1996. Throwing while looking through prisms: II. Specificity and storage of multiple gazethrow calibrations. *Brain*, 119(4), 1199–1211.
- Pohl, R. F. 2017. *Cognitive Illusions: Intriguing Phenomena in Thinking, Judgment and Memory*. Routledge/Taylor & Francis Group.
- Rosenberg, S., Vedlitz, A., Cowman, D. F., & Zahran, S. 2010. Climate Change: a Profile of US Climate Scientists' Perspectives. *Climate Change*.
- Sosa, E. 2010. The Epistemology of Disagreement. In *Armchair Philosophy*. Princeton University Press.
- Wellman, H. M., Cross, D., & Watson, J. 2001. Meta-analysis of theory-of-mind development: The truth about false belief. *Child Development*, 72(3), 655–684.
- White, R. 2005. Epistemic Permissiveness. *Philosophical Perspectives*, 19, 445–459.
- Wimmer, H., & Perner, J. 1983. Beliefs about beliefs: Representation and constraining function of wrong beliefs in young children's understanding of deception. *Cognition*, 13(1), 103–128.