



Disagreement and epistemic improvement

Fernando Broncano-Berrocal¹ · Mona Simion²

Received: 19 September 2020 / Accepted: 25 September 2021 / Published online: 1 November 2021
© The Author(s) 2021

Abstract

This paper proposes a methodological turn for the epistemology of disagreement, away from focusing on highly idealized cases of peer disagreement and towards an increased focus on disagreement *simpliciter*. We propose and develop a normative framework for evaluating all cases of disagreement as to whether something is the case independently of their composition—i.e., independently of whether they are between peers or not. The upshot will be a norm of disagreement on which what one should do when faced with a disagreeing party is to improve the epistemic properties of one's doxastic attitude or, alternatively, hold steadfast.

Keywords Disagreement · Conciliationism · Steadfast views · Epistemic norms · Epistemic rankings · Social epistemology

1 Introduction

The epistemology of disagreement is a well-established research field in social epistemology. The promise of the literature is to give an answer to the question of *what*

This article belongs to the topical collection "New Directions in Social Epistemology", edited by Adam Carter and Christoph Kelp.

✉ Fernando Broncano-Berrocal
f.broncano-berrocal@ub.edu
<https://www.fbroncano-berrocal.com/>;
http://www.ub.edu/grc_logos/

Mona Simion
mona.simion@glasgow.ac.uk
<http://www.mona-simion.com/>;
<http://www.cogito-glasgow.com/index.html>

¹ University of Barcelona – LOGOS, Barcelona, Spain

² University of Glasgow, Glasgow, UK

is permissible¹ for one to believe in the face of disagreement (Lackey, 2013). Unfortunately, this promise has not materialized. So far, the debate has almost exclusively revolved around one very specific type of disagreement, viz. *peer disagreement*, which is an idealized (or at best rare) kind of disagreement between individuals with roughly the same evidence and cognitive capacities. In this way, the guiding question of the literature has turned out to be this one instead: What would it be rational for two acknowledged epistemic peers to believe upon discovering that they are in disagreement?²

If the original aim was to offer a principled account of the epistemically permissible response to disagreement in general, why is it that epistemologists have spent so much time focusing on a small subset of idealized disagreements? The thought is that, by starting from this idealized position, we can isolate ourselves from the noise of real-life disagreement. Subsequently, the thought goes, we can upload context—i.e. add the details of particular cases of real-life disagreement to our diagnosis of what is permissible to believe in idealized cases—and thus expand one’s view on peer disagreement to real-life disagreement.

There is no doubt that this widespread methodology has helped shed light on a number of significant questions, such as the role of higher-order evidence (Christensen, 2010; Kelly, 2010; Lasonen-Aarnio, 2014), personal information (Lackey, 2010; van Inwagen, 2010), justification (Lackey, 2010), or the question of whether the same body of evidence can rationalize opposite beliefs (Feldman, 2007; White, 2005). Crucially, though, the question of what is permissible for one to believe in the face of disagreement remains unsettled to date.

This paper proposes a methodological turn. The main aim of this paper is to propose an alternative, more comprehensive normative framework for evaluating *all* cases of disagreement as to whether something is the case independently of their composition—i.e., independently of whether they are between peers or not. As such, the ambition is to answer the chief normative question—what is permissible for one to believe in the face of disagreement as to whether *p*?—for all, including everyday cases of disagreement, between, e.g., people of varied cognitive talents, and whose doxastic attitudes enjoy substantively different epistemic properties.

The paper is structured as follows. In Sect. 2, we distinguish two forms of epistemic evaluation, drawing on a distinction between source epistemic properties and state epistemic properties on one hand, and prescriptive and evaluative epistemic norms on the other. In Sect. 3, we use these distinctions to illustrate the different symmetries and asymmetries that are present in cases of doxastic disagreement. In Sect. 4, we criticize

¹ In line with the literature, we are keeping with permissibility as pertaining to the epistemic domain. Henceforth, we will take this as read.

² The answers to this question are typically grouped into two main camps. Some epistemologists (Feldman, 2005, 2007, 2009; Elga, 2007, 2010; Christensen, 2007, 2009, 2013; Bogardus, 2009; Kornblith, 2010, 2013; Cohen, 2013; Pittard, 2015; Vavova, 2014; Matheson, 2015) argue that the discovery of peer disagreement rationally compels peers to revise their beliefs, in the sense that they should decrease their confidence in what they believe. Supporters of this answer are called *conciliationists* or *conformists*. By contrast, other epistemologists (Kelly, 2005; Bergmann, 2009; van Inwagen, 2010; Wedgwood, 2010; Weintraub, 2013; Weatherston, 2013; Decker, 2014; Titelbaum, 2015) defend that epistemic peers are at least sometimes rationally entitled to stick to their guns. Supporters of this idea are known as *steadfasters* or *nonconformists*. Finally, some views lie squarely in-between these extremes (Kelly, 2010; Lackey, 2008, 2010).

what we consider an excellent predecessor to our view (the *Knowledge Disagreement Norm*), and, in Sect. 5, we propose our own norm of disagreement (the *Epistemic Improvement Disagreement Norm*), which models the permissible doxastic response to disagreement in terms of the ranking of state epistemic properties. Finally, in Sect. 6, we show how the norm handles some standard cases of disagreement.

2 Two forms of epistemic evaluation

To introduce our methodological alternative, we first need to draw a distinction between *source epistemic properties* and the *epistemic properties of one's doxastic state* (state epistemic properties, for short). What we dub 'source epistemic properties' refer to the epistemically relevant properties of the method of belief formation: is the process of belief formation reliable? Does the subject have enough evidence for her belief? Does the subject have access to her reasons for holding the belief? Depending on one's theoretical commitments, some of these properties will be relevant to whether the belief in question is justified or not. Independently of one's theoretical commitments, however, we can assess an epistemic source for these properties. In contrast, state epistemic properties pertain to the properties of the belief itself. Some of the source properties will be inherited by the state as well: justification is a central example. Others, however, will, at least on most views,³ pertain to the state alone: 'Is the belief true or false?' is one paradigmatic question that pertains to state properties only.

In order to illustrate the normative relevance of this distinction, it will be useful to take a detour into general normativity theory. According to a widely embraced view about the nature of norms, one can distinguish between evaluative and prescriptive norms.⁴ *Prescriptive norms* are primarily about what one ought to do; they prescribe certain pieces of conduct.⁵ Paradigmatic examples of prescriptive norms include moral norms such as 'Don't steal' or 'Don't lie' but also traffic norms such as 'Drive 50 km/h within city bounds' and rules of games such as 'only move the bishop diagonally' in chess. We take it that one crucial function of prescriptive norms is to reinforce certain forms of conduct (permitted ones) and to discourage others (prohibited ones). For instance, one function of the traffic norm that requires people to drive no faster than 50 km/h within city bounds is to reinforce staying within the 50 km/h speed limit and to discourage speeding.

In contrast, *evaluative norms* regulate what it takes for a token of a particular type to be good or bad with regard to its type. Take, for instance, the norm that a good hospital is a clean hospital, that a good knife is sharp or that good driving is safe driving. Evaluative norms use 'good' in Geach's (1956) attributive sense (McHugh,

³ A notable exception is Williamson's (2000) JB = K view that collapses the distinction altogether.

⁴ See, e.g., McHugh (2012) and Simion et al. (2016).

⁵ This is not to say that they need to be action-guiding in a strongly internalist sense, that is, what Williamson (2008) dubs *perfectly operationalizable norms*: norms which are such that one is always in a position to know whether one is complying with them. On the contrary, we take it that most prescriptive norms are not thus operationalizable. Take, for instance, the norm: 'Drive 50 km/h within city bounds!' Whether one complies with this norm is not something that one is always in a position to know. If this is not immediately obvious, just imagine that your speedometer is broken.

2012, p. 22), where ‘good’ functions as a predicate modifier, rather than as a predicate in its own right. When the evaluative norm states that good knives are sharp, it merely states that knives *qua knives* are good only if they are sharp. It does not entail that good knives are good *simpliciter*,⁶ or good for some purpose or another. Similarly, it might be true that good burglars are stealthy (stealth thus being a condition on complying with the evaluative norm for burglars), but this does not entail that good burglars are good *simpliciter*.

Evaluative norms differ from prescriptive norms in that they don’t prescribe a certain piece of conduct. This is not to say that evaluative norms do not feature genuine oughts. On the contrary, there is a clear sense in which evaluative norms tell us something about how things ought to be. For instance, the aforementioned evaluative norm for hospitals tells us that hospitals ought to be clean, the norm for knives tells us that knives ought to be sharp and the norm for driving tells us that driving ought to be safe. But this does not detract from the fact that such norms still do not prescribe a certain course of conduct and, as a result, are not prescriptive norms. The main difference lies in the fact that, while prescriptive norms are ought-to-dos, evaluative norms are ought-to-bes.

Evaluative and prescriptive norms can come apart. It is entirely possible for an evaluative norm to be violated without a prescriptive norm being violated. Consider for instance, again, the following evaluative norm: good knives are sharp knives. A blunt knife will be in breach of this norm. Compatibly with that, no prescriptive norm may be violated: the knives producing factory, for instance, might have done everything by the book, with this particular blunt knife being an unlucky accident. And, conversely, it is possible to violate a prescriptive norm without violating any evaluative norm. Suppose your boss has imposed a completely pointless rule for making coffee according to which coffee powder must be scooped into the filter alternatingly with a blue and a red teaspoon. Suppose you violate this rule, e.g., because you are only using the red spoon. You have broken a prescriptive rule for coffee making at your office. However, it seems plausible that you need not have violated some evaluative norm as well. In particular, the coffee you are making may still be good coffee, the way you are making it may still be a good way of making coffee and you may still be a good barista.

While evaluative norms thus differ from prescriptive norms, the two may still be (and often are) related. In particular, prescriptive norms often enough derive from evaluative norms.⁷ They serve to ensure that an evaluative norm is likely enough complied with. For instance, prescriptive norms of driving, such as the norm ‘Drive no more than 50 km/h within city bounds’, serve to ensure that the evaluative norm of driving—according to which good driving is safe driving—is likely enough complied with. In this way, evaluative norms often come first and prescriptive norms are in their service.

⁶ Note that Geach (1956) argues that there’s nothing more to good *simpliciter* than attributive goodness. We don’t take any stance on the issue here.

⁷ Deontologists will disagree: these philosophers take the deontic to explain the axiological (see, e.g., fitting attitude accounts of value, *à la* Rabinowicz & Ronnow-Rasmussen, 2004). We will set aside this complication for the purposes of this paper. See Simion (2021) for discussion.

To go back to epistemology: in the case of belief, the evaluative and prescriptive are widely taken to come apart also. How this will work will depend on the specifics of the views in question, but here is how this works at a general level: prescriptive norms will govern the methods of belief formation, while evaluative norms will govern the output of the formation exercise—they will tell us what a good belief is. Suppose that, in line with tradition, one takes the answer to the question ‘How should one go about forming beliefs?’ to be different from the answer to the question ‘What is a good belief?’. Suppose, further, that one champions one of the (arguably) main views on the aim of belief in the literature: according to the traditional view, truth is the aim of belief—when we form beliefs, the aim of this endeavor is to generate true beliefs. One competing view takes the relevant aim to be knowledge: mere belief, it is argued, is botched knowledge; good belief is knowledgeable belief.

It is easy to see that, whatever the details of the particular views, one can end up with a view on which the prescriptive and the evaluative norms of belief come apart. In particular, it is possible to violate the evaluative norm of belief whilst complying with the prescriptive norm. For illustration: if you are an evidentialist internalist and you like the truth goal of belief, on your view, truth will be the evaluative norm of belief, while the prescriptive norm will require you to accord your beliefs to your evidence. If you are a process reliabilist, you will, again, take good belief (belief meeting the evaluative norm) to be true belief and the prescriptive norm will require one to produce beliefs via reliable processes. And so on.

With all this in mind, we want to go back to our distinction between source epistemic properties and state epistemic properties. The former will be governed by prescriptive norms, tracking the epistemic permissibility of the method of belief formation in question. Questions like ‘What epistemic properties should one’s method of belief formation have?’; ‘How should one respond to available evidence?’, and so on will be answered by the prescriptive epistemic norm. In contrast, questions pertaining to state epistemic properties—such as ‘What properties should one’s belief have?’ or ‘What properties should one’s credence have?’—will be governed by the evaluative norm of belief.

Accordingly, for any doxastic attitude a subject hosts, we can distinguish *two forms of epistemic evaluation*: (1) the evaluation of source epistemic properties—by how well they meet the prescriptive norm one favours—and (2) the evaluation of the state epistemic properties—by how well they meet the evaluative norm one favours. Just like prescriptive and evaluative norms, the two forms of epistemic evaluation will plausibly often be related, in that (2) often depends on (1).

Different sets of source properties, as well as different sets of state properties, can be compared: they can be better or worse. We compare the sets of source properties for two or more agents with respect to their doxastic attitude towards a proposition by fixing one or more parameters—those central to our preferred epistemological theory—and by then seeing how well the sources fare with respect to them, i.e., by ranking them with regard to closeness to the set parameter(s). For instance, we may say that *A* is doing better than *B* when it comes to source properties because *A* has a more reliable belief-forming method, or better evidence than *B*. In general, the choice of the relevant parameters is down to one’s epistemological inclinations.

In the same vein, we can also compare (i.e., regard as better or worse) the sets of epistemic properties enjoyed by two doxastic attitudes about a proposition (e.g., beliefs, credences). To do so, once again, we refer to some ranking of sets of epistemic properties *vis-à-vis* one or more set parameters that are central to our preferred epistemological theory, and regard as epistemically better the doxastic attitude that instantiates the best set of properties on this ranking. By way of illustration, a popular way to rank state properties is this: being known is better than being merely true which is in turn better than being false. Accordingly, to take a simple toy case, if *A* believes that *p* and *B* believes that-not *p*, and *A*'s belief amounts to knowledge, whereas *B*'s belief is false, then the set of state properties of *A*'s belief is better than the set of state properties of *B*'s belief, according to this ranking. As in the case of fixing the relevant parameters for evaluating source properties, how to rank sets of state properties depends on one's epistemological sympathies. Some rankings should be mostly uncontroversial, such as the previous ranking. However, for some pairings it is unclear what set of state properties should rank higher. It is in these cases that theoretical considerations become relevant and, consequently, where controversy can be expected. For example, does unjustified true belief rank higher than justified false belief? Plausibly, this question cannot be answered drawing only on pre-theoretical considerations. The important point for our purposes is simply that sets of state properties (just as sets of source properties) can be ranked with regard to closeness to one's favourite parameter, and thus that they can be ranked in different ways.

As we will see, this will provide the theoretical ground for a new way to understand the epistemology of disagreement. Before that, and to see how the two forms of epistemic evaluation distinguished here bear on the epistemology of disagreement, let us analyze the different symmetries and asymmetries that can be found in doxastic disagreement.

3 Symmetries and asymmetries in doxastic disagreement: a methodological turn

Consider a disagreement in which *A* and *B* disagree about whether *p* is the case, where *A* believes *p* and *B* believes not-*p*. One way to evaluate their disagreement is to compare how good *A*'s and *B*'s respective sets of source epistemic properties are with respect to their doxastic attitudes, and especially (as we have seen) how well they respectively fare with respect to certain parameters, such as, e.g., the evidence they have in favor of or against *p*/not-*p*, or the reliability of their belief-forming methods.

Suppose that we agree on what parameters are relevant to make an epistemic assessment of their respective sets of source properties. There are two general possibilities: (1) that *A* and *B* do equally well with respect to all the chosen parameters (e.g., that they have the same evidence, that their belief-forming methods are equally reliable, and so on); (2) that they do not.

Peer disagreement, as generally conceived in the literature, is disagreement of the first type. In other words, peer disagreement is the kind of disagreement that takes place between two or more agents who have equally good sets of source epistemic

properties with respect to their doxastic attitudes. By contrast, ordinary disagreement is typically disagreement of the second type.

Crucially, though, evaluating the distribution of source epistemic properties is not the only way to approach the epistemology of disagreement, nor the only way to conceive of the question of what is permissible for one to believe in the face of disagreement (recall: this is the key question that the literature promises to answer). As an alternative, one can focus on the distribution of state epistemic properties. Or, put differently, one can take the permissibility at stake to map on to evaluative normativity rather than to prescriptive normativity.

The main difference is that, while accounting for the rational response to disagreement in terms of the distribution of source properties inevitably divides the cases into two sets—i.e., cases of sameness in quality of source properties and cases of asymmetrical quality of source properties—, modelling this response in terms of the distribution of state epistemic properties does not make any such division: all cases fall into the same category. Indeed, by focusing on the distribution of source epistemic properties alone, we submit, one leaves out of the equation the most fundamental asymmetry that is present in *all* disagreements about whether something is the case:

The fundamental asymmetry: The doxastic attitudes of the disagreeing parties as to whether p is the case never have identical sets of epistemic properties.⁸

Whereas A knows that p , B holds a false belief that not- p ; whereas A suspends judgment on p , B holds a true/false belief about p , and so on. Of course, what we *never* see in a disagreement as to whether p is that one party knows or holds a true belief that p and that the other party knows or holds a true belief that not- p , or that both parties suspend judgment on p /not- p (those are simply not disagreements).

We believe it is time to try an alternative methodology for the epistemology of disagreement. In a nutshell, we think that the alternative is to take into account what has been left out of the equation so far and what defines, in our opinion, the subject matter: the fact that the doxastic attitudes of the disagreeing parties never have identical sets of epistemic properties. This is what we have called the *fundamental asymmetry*, and is at the core of the *methodological turn* we propose for the epistemology of disagreement.

In particular, by accounting for the epistemology of disagreement in terms of what all disagreements as to whether something is the case have in common, i.e., in terms of this fundamental asymmetry, we can begin to map all possible cases independently of whether they are instances of idealized, peer disagreement or else of real-life, asymmetrical disagreement. Indeed, that a given case is a case of peer or non-peer

⁸ This assumes a view of disagreement that, to some extent, discounts cases in which you and I have different credences in p . Some people—although others disagree—think (some of) these are genuine cases of disagreement. The way to accommodate the intuition that these are cases in which we genuinely disagree is to work on the content: we disagree about how confident one should be about p . Similarly, we take it that it is plausible that in cases in which one party suspends on p (e.g., that God exists) and the other believes that p , there is disagreement about what attitude the evidence supports. In a nutshell, then, we take these disagreements to be genuine disagreements, but to be placed at the higher-order level—i.e., to concern the appropriate doxastic attitude to have rather than the content thereof.

disagreement is orthogonal to the distribution of state epistemic properties, but rather depends on the distribution of source epistemic properties.⁹

Some caveats are in order to avoid misunderstandings. First, by arguing for the existence of a fundamental asymmetry shared by all cases of disagreement as to whether p , we are not claiming that the doxastic attitudes of the disagreeing parties can't share *some* epistemic properties, but only that, necessarily, they can't *all* coincide. For if the two sets of state epistemic properties are identical, there is no disagreement as to whether p .

Second, what we call the 'set of state epistemic properties' is not value-charged, but it rather pertains to the overall set of properties enjoyed by the doxastic attitude in question. That is, one might have different intuitions (or theories) about how such state properties are or should be *ranked*, but that is orthogonal to the *existence* of the fundamental asymmetry. For the asymmetry follows from the very fact that A and B are in disagreement as to whether p and hence precedes any epistemic evaluation thereof. Note that in a disagreement as to *whether* p , the asymmetry in the two relevant sets of state epistemic properties can be at least explained by appealing to the fact that A 's belief in p is true whereas B 's in not- p is false (or vice versa). But one may also appeal to differences in accuracy, justification, or in other epistemic properties. In brief, to characterize the fundamental asymmetry we only need to point to one or several differences in state epistemic properties while staying neutral on how to evaluate or rank them.

Third, and relatedly, our methodology is neutral on and, indeed, compatible with both internalist and externalist conceptions of state epistemic properties. Take epistemic justification. Our methodological approach in terms of the fundamental asymmetry stays neutral on whether justified true belief should, for instance, rank as high as justified false belief because justification is conceived exclusively in terms of internal factors (e.g., factors that are reflectively accessible to the agent). The question as to whether only internal stuff matters simply does not pertain to what state properties a belief instantiates, but it rather pertains to how to *rank* sets of such properties: indeed, one may have a view on which internal goodness is all that matters, and, as such, justified true beliefs and justified false beliefs will enjoy equally good sets of epistemic properties on this view. But they will still instantiate *different* sets of epistemic properties nevertheless, and that's all we need to commit to for the kind of methodological turn we're proposing.

This will be an important aspect of our account, which is intended to be broad and so *non-committal* with respect to the more specific norms of disagreement that may arise from applying different rankings of state properties resulting from different epistemological theories and sympathies. For this reason, our account of disagreement should not be read as implicitly endorsing a kind of epistemic relativism, but rather a lack of commitment to specific epistemological theories. We take this to be an advantage of the view.¹⁰

⁹ For simplicity, we will focus on cases of disagreement with two parties. But the resulting view can be easily extended to cases in which more than two parties are in disagreement, as well as to cases of group disagreement. This is, we submit, another advantage of the alternative methodology we propose.

¹⁰ Many thanks to an anonymous referee for urging us to clarify this point, and, in general, for tremendous help with this paper.

With these caveats in place, in the next section, we will take a close look at what we think is a predecessor of the view that we will propose, insofar as it accounts for what is permissible for one to believe in the face of disagreement in terms of a specific way to rank state epistemic properties. It is the *Knowledge Disagreement Norm* (KDN), by John Hawthorne and Amia Srinivasan (2013). The hope is that, by critically assessing their view, we will take further steps in setting the ground for ours.

4 A promising predecessor: the knowledge disagreement norm

According to John Hawthorne and Amia Srinivasan (2013) (henceforth H&S), the thing to do when faced with disagreement as to whether something is the case is, as with all other epistemic activities we engage in, to aim to gain knowledge. Furthermore, according to them, “[f]rom the perspective of a knowledge-centric epistemology—that is, an epistemology that takes the most central goal of our epistemic activity to be knowledge—it is natural to rank outcomes with knowledge over outcomes of withholding belief, which are in turn ranked over outcomes of knowledge-less belief” (H&S 2013: 11). As such, they propose the following norm of disagreement:

The Knowledge Disagreement Norm (KDN): In a situation where *A* believes that *p* and *B* believes that not-*p*:

- (i) *A* ought to trust *B* and believe that not-*p* if and only if were *A* to trust *B*, this would result in *A*’s knowing not-*p*,
- (ii) *A* ought to dismiss *B* and continue to believe that *p* if and only if, were *A* to stick to her guns, this would result in *A*’s knowing *p*, and
- (iii) in all other cases, *S* ought to suspend judgment about whether *p*.

Again, according to H&S, KDN is readily borne out by epistemologies that take knowledge to be of central epistemic concern (call this thesis *Knowledge-First Support*). In particular, one claim that such epistemologies typically endorse is:

The Knowledge Norm of Belief: One’s belief that *p* is epistemically permissible only if one knows that *p*.

According to this picture, then, the norm of belief in disagreement situations coincides with the general norm of belief.

There are several important benefits to H&S’ proposal. First, KDN transcends the peer disagreement debate and talks about the right thing to do in general when disagreeing. In addition to this, KDN brings the disagreement debate in line with the latest developments in general epistemic normativity. On their view, the question as to what we are permitted to believe in cases of disagreement is just a species of the more general question: ‘What is permissible belief?’. Furthermore, if H&S are right, and KDN simply “falls out” of a knowledge-centric epistemology, it will stand to inherit all the theoretical upsides of the latter: independent plausibility, elegance, and straightforward solutions to a bunch of puzzling epistemological issues, such as, to name but a few, Gettier cases, epistemic Frankfurt cases, the lottery paradox, the paradoxical soundingness of Moorean sentences or lack of justification in clairvoyance cases.

All this being said, there are two concerns about KDN that we want to set aside.

The first relates to the possibility of defeat: intuitively, at least in some cases, when involved in an everyday disagreement with someone the epistemic credentials of whom I have no reason to doubt, I should lower my confidence in the target proposition: testimony is a *bona fide* source of defeat. Of course, given this, the possibility of acquiring knowledge will not be present in the vast majority of disagreement cases, so it would seem that KDN would recommend generalized withholding, which sounds extreme. However, traditionally, many knowledge-first epistemologists (H&S included) reject knowledge defeat (see, e.g., Lasonen-Aarnio, 2014). According to these people, one is entitled to hold on to one's belief in the face of evidence to the contrary if one knows, and intuitions of impermissible dogmatism are explained away via error theory: we feel these folks are impermissibly dogmatic because they are epistemically blameworthy ignoring evidence, not because they epistemically impermissibly do so. We are ourselves skeptical about this move. However, for the purposes of this paper, we will take it on board and assume it works.

Another concern about KDN that readily comes to mind, as with knowledge norms in general, is its little potential for being intuitively satisfactory. Note that KDN looks nothing like the norm we were after when asking what we should do when faced with disagreement. After all, often enough, one is hardly in a position to know who is the knowledgeable party. Of course, if I were to know that you are the one who knows, and, therefore, that you are right, I should stop disagreeing. The point, however, is that in disagreement situations, that is one piece of information that will typically be missing. As such, one might worry, KDN suffers from *lack of operationalizability*.

H&S acknowledge as much. They note, however, that, when it comes to normativity, lack of transparency looms larger than commonly assumed, i.e., much more internalistic norms than KDN will suffer from it also. After all, we do not only have imperfect access to the outside world, but also to our own minds. If that is the case, H&S argue, the prospects of formulating a perfectly luminous, and therefore intuitively satisfactory epistemic norm, no matter how far internalist we are willing to go, look rather dim. We strongly agree with H&S on this count, and we will assume the correctness of their defense throughout this paper.

To sum up what we've seen so far, then: at first glance, KDN looks good; it comes with important theoretical upsides and, furthermore, the price one needs to pay for endorsing it (i.e., non-luminosity) is, at a closer look, not that high after all, for being shared by all competitors on the market.

However, we have worries for KDN sourced in reasons other than those discussed by H&S. In particular, we think that *Knowledge-First Support* is false. Against H&S, we will argue that KDN does not trivially “fall out” of a knowledge-centric epistemology. Furthermore, we will argue that, quite to the contrary, it does not even snugly fit within such a picture. This is because, if knowledge is, indeed, the main telos of all of our epistemic endeavors, KDN is not merely non-luminous: it is *uninformative*. Furthermore, as we will try to show, KDN, in conjunction with a knowledge-centric picture gives rise to implausible value-theoretic results.

Let us start by going back and unpacking the rationale behind KDN. Recall that, according to H&S, the following holds:

The Knowledge Goal Thesis: Knowledge is the telos of our epistemic activity.

Furthermore, they take it that, in view of the fact that knowledge is our main epistemic goal, it follows that we should (1) only believe what we know, and (2) otherwise withhold. Let us try to state this assumption more precisely. It looks as though H&S stand behind the following two theses:

The Knowledge Goal—Knowledge Norm Link: If the *Knowledge Goal Thesis* is true, one should only believe what one knows.

The Knowledge Ranking Thesis: If the *Knowledge Goal Thesis* is true, outcomes with knowledge are better than outcomes of withholding belief, which are better than outcomes with knowledge-less belief.

What we will argue next is that both the *Knowledge Goal—Knowledge Norm Link* and the *Knowledge Ranking Thesis* are problematic.

4.1 Against the knowledge goal: knowledge norm link

There are reasons to doubt that the *Knowledge Goal—Knowledge Norm Link* holds. To see this, think of other norms we are familiar with. Consider traffic norms. Plausibly, they are meant to promote the aim of producing safe traffic. Now, note that norms meant to make it likely that the goal of safe traffic is reached have *informative* content, i.e., content that tells us how to go about reaching the goal in informative ways: ‘Drive at most 50 km/h within city bounds!’, ‘Stop at the red light!’, and so on. Similarly, consider social norms. Suppose that you want to become a biologist. Again, there are going to be norms that inform you how to go about it in order to make it likely that you do become a biologist: ‘Go to biology school!’, ‘Study hard!’, and so on. Or consider norms of chemistry. Suppose that you want to produce an antibiotic. The norms associated with this goal are going to tell you what ingredients to mix (and under what conditions) in order to produce an antibiotic. It would be less than informative if all these norms would *share content with the goal*. For example, if the norm serving the goal of your becoming a biologist merely said ‘Become a biologist!’, or the norms regulating antibiotics production merely said ‘Produce antibiotics!’. Similarly, it would be less than informative if the only driving norm we would have to guide us towards our goal of safe traffic were ‘Drive safely!’.

Now, unless we have good reason to think that epistemic normativity is special in this respect, we should expect to get a similar picture for belief: the norms should not share content with the goal. They should be such that they inform us how to go about belief formation, and that compliance with them makes it likely that you end up with knowledgeable beliefs. A norm that tells us to only believe what we know does not seem to guide us much with regard to how to go about reaching our goal of producing knowledgeable beliefs.

A similar argument, this time in the literature on the *knowledge norm of assertion*, is due to McKinnon (2015). The key thought is that, on many versions of this norm of assertion,

[...] the goal and the norm of the practice are satisfied by the same thing: knowledge. [...] However, [...] this is a serious mistake: the norms and goals of a practice must be kept distinct, because norms are what it means to “properly aim” at a goal. (McKinnon, 2015: 170)

To go back to the *Knowledge Goal—Knowledge Norm Link*: one should distinguish between *perfect operationalizability* of a norm, which demands full luminosity and is therefore, as H&S convincingly argue, an unreasonable expectation, and *informativeness*, which seems like a fair desideratum to impose on any satisfactory normative account. Again, we strongly agree: norms in general are not perfectly operationalizable, due to featuring non-luminous conditions for satisfaction—e.g., I am not always in a position to know whether I am, indeed, driving 50 km/h. However, norms in general tend to be informative towards reaching the goal they observe, that is, they tend *not* to share content with the goal.

It turns out, then, that even if we grant H&S the generalized non-luminosity claim, there is a more serious worry looming in the background: if knowledge is the goal, a knowledge norm is not very informative towards reaching it. As such, we have reason to suspect that the *Knowledge Goal—Knowledge Norm Link* is, to say the least, not as straightforward as H&S take it to be. But if that is the case, KDN risks to remain unmotivated.

4.2 Against the knowledge ranking thesis

Let us now have a closer look at the *Knowledge Ranking Thesis*. We are ready to grant H&S that, if knowledge is our epistemic goal, outcomes with knowledge will be epistemically better than outcomes without knowledge. What seems a bit trickier, however, is to understand why knowledge-centric epistemology needs to rank withholding belief above *all* varieties of non-knowledgeable belief. After all, the latter can come in pretty sweet shapes, epistemically. Think, for instance, of true belief which enjoys quite some degree of justification, just not enough to qualify as knowledge. Why is it that, when motivated by a knowledge goal, one should abandon this belief in favour of withholding? After all, it looks like I am closer to knowledge if I hold on to my belief than if I abandon it. Furthermore, think of the disagreeing party. This fellow is faced with a nice opportunity to move from a false belief to a true one, and get some justification for it too. According to KDN, however, he should miss out on this opportunity.

To see why the KDN claim is further problematic, note that it implies a fairly unusual view on epistemic value. After all, starting with Plato’s *Meno*, epistemologists have struggled to answer the question of why knowledge is more valuable than epistemic standings that fall short of knowledge. And the answer has hardly proved to be straightforward. But even if we leave this aside, there is still a long way left to go for justifying H&S’s *Knowledge Ranking Thesis*. After all, according to their ranking, even withholding belief is more valuable than any epistemic standing that falls short of knowledge. Apart from the fact that the thesis sounds less than plausible to us, it surely affords quite a bit of argument, and it hardly follows naturally from a knowledge-centric account, as H&S would have it. After all, a knowledge-centric

account need not hold (and, arguably, shouldn't) that truth is totally worthless. Rather, a more plausible picture would be one where, roughly, knowledge is 'gold' and truth is 'silver.' But if that is the case, it is unclear why one should rather suspend judgment than hold a true belief.

One way for H&S to try to defend the *Knowledge Ranking Thesis* would be via the *Knowledge Goal—Knowledge Norm Link*, i.e., the thesis that, if knowledge is the goal of inquiry, one should only believe what one knows. After all, one might think that if the *Knowledge Goal—Knowledge Norm Link* is true, we have an answer to the question of why is it that withholding belief is better than any epistemic standing that falls short of knowledge. The answer is pretty straightforward: because, by withholding, one avoids breaking the norm of belief.

Unfortunately for H&S, this move will not work so easily either. First, more is needed to move from the *Knowledge Goal—Knowledge Norm Link* to the *Knowledge Ranking Thesis*. In other words, H&S need an argument to the effect that not being in breach of the norm of belief is more valuable than being in the possession of truth. This, again, affords quite a bit of argument. Furthermore, as we have already seen above, there are reasons to be suspicious of the *Knowledge Goal—Knowledge Norm Link* itself.

Another way for H&S to try to avoid the counterintuitive value-theoretic results would be to fall back on contrary-to-duty normativity (Chisholm, 1963). In many cases our moral obligations depend essentially upon our moral failings. For example, a person might have a moral obligation to apologize for some misbehavior—yet obviously he would not be apologizing if he hadn't already broken a norm. It is easy to see how one might think advocates of KDN could invoke the contrary-to-duty ought to pick up the intuitions that one should do the best one can given that one cannot (or will not) do what one ought to do. As such, the thought would go, KDN can accommodate the thought that, short of knowledge, one should believe, e.g., truly and justifiably, and, short of that, one should believe truly, and, in turn, short of that, one should withhold.

Unfortunately, this move will not work for KDN either, in its current formulation. Since withholding is an alternative to knowing built in the content of KDN, it will rank higher than any epistemic standing that is featured in the content of the corresponding contrary-to-duty imperative. To see this, say that the contrary-to-duty imperative to KDN would read something like: when in breach of KDN, one should believe truly. It is easy to see that, on this picture, one should try to gain knowledge, if that doesn't work one should withhold, and only if one fails to do so, go ahead and make an attempt at acquiring a true belief. In this, again, withholding is ranked as more epistemically valuable than true belief.

What would need to happen for this fallback to become available to the KDN champion would be to reformulate KDN as a simple knowledge norm. In the next section, we put forth a norm of disagreement that is compatible with a view along these lines: if taken in conjunction with a knowledge goal of belief, it will deliver the intuitively right value-theoretic ranking we wanted from KDN.

5 The epistemic improvement approach to the epistemology of disagreement

We believe that the problems facing KDN reveal an important fact about our epistemic life, where disagreement makes no exception to the rule: inquiry is not an all or nothing affair. As our critical discussion of H&S's proposal has made transparent, if the epistemically permissible response to disagreement were the one that KDN predicts, one would lose the opportunity to form beliefs with many desirable epistemic properties. Recall the case of someone having a true belief that p which enjoys quite some degree of justification—just not enough to qualify as knowledge—having a disagreement with someone who holds an unjustified false belief that not- p . On the 'all-or-nothing' KDN model, the party who is right is supposed to switch to withholding in the light of disagreeing. One might want to avoid this result. Furthermore, it certainly seems as though this is a good moment for the disagreeing party to abandon her unjustified false belief and adopt the belief that p .

Here is one thought we find plausible: inquiry is a goal-directed epistemic practice, and disagreement is a move in inquiry. Now, it's a fact that there is disagreement about the precise goal of inquiry: Is it truth? Knowledge? Justified belief? However, one thing that should be hardly controversial is that moves in practices aim—in either a direct or an indirect way—to fulfil the goal of the practice. Diabetological consults are moves in the practice of medicine, and they aim directly at fulfilling the goal of the practice of medicine: curing diseases. In turn, performing blood sugar tests aims (directly) at informing the diabetologist as to how well the patient is doing, which, in turn, aims directly at her diagnosing the patient correctly, and, further, at curing her disease. In this, blood sugar tests aim indirectly at the goal of the practice of medicine: curing diseases. They aim at making progress towards it. Similarly, baking cakes is a move in the practice of cooking, and it aims directly at fulfilling the aim of the practice: producing tasty, nourishing food. My getting the flour off the shelf aims indirectly at the general aim of the practice by aiming directly at adding flour to the cake mix. It aims at making progress towards producing tasty, nourishing food. And so on.

If that is the case, though, in virtue of being a move in our practice of inquiry, disagreement will aim (directly or indirectly) to fulfill the goal of the practice: it will aim, then, at a minimum, at making progress towards it.

In the light of all this, it looks as though if we are interested in getting anything out of disagreeing with each other as to whether something is the case, it is in *epistemically improving* towards our epistemic goals, whatever they may be, via adopting the doxastic attitude with the best set of epistemic properties available, or else retain one's current doxastic attitude if no upgrade in state epistemic properties can be obtained.

By way of illustration, suppose that you are in the business of finding out whether there is milk in the fridge. Given your current epistemic situation, you only have two options: (1) open the fridge and see whether there is a bottle of milk; (2) ask your 3-year-old son whether there is milk in the fridge. The best you can do in this situation, epistemically, is to open the fridge. Now imagine a different situation in which, unbeknownst to you, your 3-year-old son has replaced the milk in the fridge with fake milk. You have the same options as before: (1) check whether there is milk in the fridge and (2) ask your son. The best you can do now is the latter. Why? Because

if you checked whether there is milk in the fridge by yourself, you would be deceived and form a false belief. By contrast, if you asked your son, being the terrible liar that he is, you would notice that he is lying and therefore come to know that there is no milk in the fridge.

A disagreement is a special kind of situation in which the options for belief formation, maintenance and update are mainly two: (1) conciliating (i.e., revising one's doxastic attitude in the direction of the doxastic attitude held by the disagreeing party) or (2) holding steadfast. With this in place, one can anticipate the *key ideas* of our proposal, which can be *roughly* summarized as follows, in the wait of more precise formulation:

The Epistemic Improvement Approach to Disagreement (EI): In a disagreement as to whether p is the case, one should improve the epistemic properties of one's doxastic attitude about p as much as possible.

To be sure, the above is not yet a precise norm of disagreement, but just a rough outline of the approach we favour: instead of accounting for the epistemology of disagreement in terms of the distribution of source epistemic properties, EI models it in terms of the *distribution of state epistemic properties*. Suppose that A believes that childhood vaccination is safe, whereas B denies it. EI's way to analyze the epistemic normativity that governs A and B 's disagreement is to focus on what epistemic properties their respective beliefs instantiate, e.g., A may know or justifiably believe that childhood vaccination is safe, whereas B unjustifiably believes the denial of this proposition. If that is the case, the epistemically permissible response to this disagreement would be, according to EI, for B to conciliate, as by doing so she would gain (at least) a true belief, and for A to stick to her guns, as by doing so she would retain the better state epistemic properties of her belief.

With the core ideas of our proposal in place, we proceed to put forward a more specific and precise norm of disagreement.

5.1 First stab: doxastic attitudes, epistemic properties, and rankings

EI models the rational response to disagreement as to whether p in terms of the distribution of state epistemic properties, but state epistemic properties can be ranked in different ways. The ambition is for the norm to be compatible with *whatever* way to rank state epistemic properties one prefers. This doesn't mean that the norm will be fully dependent on one's epistemological sympathies: insofar as everyone will accept certain uncontroversial ways to rank state epistemic properties, there should be no controversy concerning the relevant direction of epistemic improvement in disagreements featuring them (actually, we submit, most disagreements).

In particular, the direction of epistemic improvement should be uncontroversial in disagreements featuring, for instance, pairings of true belief *simpliciter* versus false belief *simpliciter*, pairings of unjustified false belief versus justified true belief, and pairings in which one's doxastic attitude enjoys a set of epistemic properties that entails the set of epistemic properties of the disagreeing party's doxastic attitude (e.g., knowledge versus justified belief). Of course, for some other pairings (such

as disagreements involving unjustified true beliefs versus justified false beliefs), we should expect some controversy, so long as the question of which one ranks higher cannot be answered in the absence of theoretical considerations. This means that, depending on one's epistemic normative commitments, different answers might follow and, as we have seen, will afford argument. We want our epistemic improvement norm to be compatible with *any* such commitments.

Moreover, we want our norm to be compatible with different views on what the *central goal* of our epistemic lives is (such as truth or knowledge) and corresponding views about what it takes to get closer to such a goal. With that in mind, here is, then, our first stab at the norm of disagreement:

The Epistemic Improvement Disagreement Norm₁(EIDN₁): Given a value ranking R of sets of state epistemic properties mapped with respect to proximity to epistemic goal G , for all cases in which A and B disagree about whether p —where A has a doxastic attitude with content p with a set of epistemic properties E_1 and B has a doxastic attitude with content $\text{not-}p$ with a set of epistemic properties E_2 :

- (i) If E_1 is lesser than E_2 on R , then A should conciliate.
- (ii) If E_1 is equally good¹¹ or better than E_2 on R , then A should hold steadfast.

We submit that EIDN₁ meets the neutrality desideratum outlined above with flying colours. G stands for the state epistemic property taken to constitute the chief epistemic goal in the framework at stake. It is easy to see that EIDN₁ is compatible with any view on the matter. For illustration, if one thinks that the chief epistemic goal is knowledge, the ranking at work in EIDN₁ will map on to closeness to knowledge. Furthermore, EIDN₁ is compatible with different rankings concerning closeness to knowledge.

5.2 Second stab: prima facie and basing conditions

Here are two problems that immediately arise for EIDN₁. First, in a disagreement between A and B , it could be the case that A conciliates, i.e., that A adopts B 's doxastic attitude, and that this results in (in the sense that it contributes to) an upgrade in state epistemic properties for A , while at the same time other factors also partially contribute to the upgrade (e.g., in parallel, A finds further evidence in favour of B 's doxastic attitude). This is precisely the kind of noise that we want to isolate from: after all, what

¹¹ For example, a purely internalist ranking may rank a justified false belief and a justified true belief equally high. If any such a ranking were plausible (we are not, ourselves, big fans of it, but we do want our norm to stay as neutral as possible on this debate), our account predicts that one should hold steadfast in cases of this sort. Why not be permissive instead? Why not allow both for conciliatory and steadfast responses to these cases? Here are a few reasons for this. First, changing one's doxastic attitude for an equally valuable one seems like a waste of intellectual energy, and thereby practically irrational. Second, on a value-first view of reasons, you would thereby be acting without reason, since there's no increase in value to be had. Third, and most importantly, conciliation may lead to epistemic problems such as the following. Consider a case in which you have a justified true belief that is very coherent with your entire web of beliefs. Say that you abandon it for a justified false belief, as per above. Now you will get massive conflict in your web of beliefs, and something's got to go. Should you decide to update your entire web of (true) beliefs to match this new one, that would be an epistemic catastrophe. So it'd be better not to conciliate in the first place.

we are interested in providing is a *prima facie* norm of disagreement—i.e., a norm that states what one should do in virtue of being faced with the disagreement alone, independently of any other evidential considerations. Accordingly, EIDN₁ requires a *prima facie* proviso.

Relatedly, on EIDN₁, it could be the case that one complies with the norm simply by, for instance, adopting the position of the disagreeing party for no good reason whatsoever, rather than *in virtue of* the disagreement itself. Take someone who, after a disagreement about whether *p*, moves from a belief that *p* to a belief that not-*p* independently of the disagreement itself as follows: as soon as she hears the disagreeing party assert ‘not-*p*,’ a strong wish that not-*p* be the case arises in her, and that’s what generates the corresponding belief. Similar considerations apply to holding steadfast: someone might stick to their guns motivated by practical reasons that have nothing to do with the disagreement itself.

What the epistemic improvement disagreement norm should say is that it is *in virtue of* the disagreeing party’s testimony that the move to conciliate/hold steadfast happens. Here is our second stab at this:

The Epistemic Improvement Disagreement Norm₂(EIDN₂): Given a value ranking *R* of sets of state epistemic properties mapped with respect to proximity to epistemic goal *G*, for all cases in which *A* and *B* disagree about whether *p*—where *A* has a doxastic attitude with content *p* with a set of epistemic properties *E*₁ and *B* has a doxastic attitude with content not-*p* with a set of epistemic properties *E*₂:

- (i) If *E*₁ is lesser than *E*₂ on *R*, then *A* *prima facie* should conciliate *in virtue of* her disagreement with *B*.
- (ii) If *E*₁ is equally good or better than *E*₂ on *R*, then *A* *prima facie* should hold steadfast *in virtue of* her disagreement with *B*.

5.3 Third stab: the epistemic impact of disagreement

As formulated so far, the epistemic improvement disagreement norm takes as inputs the state properties present *before* the disagreement occurs, and outputs recommendations as to how to proceed to secure epistemic improvement. This, however, misses one important datum: the impact of the disagreement itself on the distribution of state epistemic properties.¹²

To see this, note that the set of properties of one’s doxastic attitude post-disagreement may be very different from the properties one starts with, in virtue of epistemic defeat (loss/partial loss of justification) or doxastic defeat (loss of belief/credence level). This in itself need not generate differences between the ranking before disagreement and the ranking post-disagreement. However, crucially, it *might* do so. For illustration, contrast the following two cases of defeat:

Cold 1

¹² Many thanks to an anonymous referee for pressing us on this.

I truly and justifiably believe I have a cold: I have all the symptoms, and I'm good at recognizing them. My friend Paul, however, disagrees: he thinks that it's not a cold, that my cough sounds like it's pneumonia, and that I should go to the hospital. I keep thinking it's a cold, but I'm less confident than before.

In *Cold 1*, defeat from disagreement with my friend downgrades my initially fully justified belief that I have a cold to a somewhat justified credence that I do. However, (at least, but not only, on the assumption of a truth-sensitive ranking R), my doxastic attitude remains strongly superior to Paul's, just as we started off before disagreement. Consider now this other case:

Cold 2

I truly and justifiably believe I have a cold: I have all the symptoms, and I'm good at recognizing them. My doctor, however, disagrees: she thinks that it's not a cold, that my cough sounds like it's pneumonia, and that I should go to the hospital. I abandon my belief that it's a cold, suspend judgement, and head to the hospital.

In this disagreement, intuitively, the epistemically permissible thing for me to do in the face of expert testimony is to conciliate, i.e., in the sense that I should update my doxastic attitude in the direction of the expert's. Indeed, my suspension on the issue seems justified. Note, though, that on many rankings R , I will have thereby lost my epistemic advantage over my doctor: while I started off with a fully justified true belief—and my doctor hosted a justified false belief—, I am now in justified suspension. Is my justified suspension still epistemically superior to my doctor's justified belief? Not clear: some rankings might predict it's not. Indeed, for instance, a purely internalist, justification-centric ranking would easily predict they are equally valuable.

If all this is so—i.e. if rankings themselves can change post-disagreement—, it looks as though, if the norm is to take the epistemic impact of disagreement into account, it should take *post-disagreement* state epistemic properties as inputs, rather than *pre-disagreement* state properties. Here it goes:

The Epistemic Improvement Disagreement Norm₃ (EIDN₃): Given a value ranking R of sets of state epistemic properties mapped with respect to proximity to epistemic goal G , for all cases in which A and B disagree about whether p —where, *after the discovery of the disagreement*, A has a doxastic attitude with content p with a set of epistemic properties E_1 and B has a doxastic attitude with content not- p with a set of epistemic properties E_2 :

- (i) If E_1 is lesser than E_2 on R , then A *prima facie* should conciliate in virtue of her disagreement with B .
- (ii) If E_1 is equally good or better than E_2 on R , then A *prima facie* should hold steadfast in virtue of her disagreement with B .

5.4 Final account: conciliation and epistemic improvement

Suppose that A conciliates, as EIDN₃ demands, because her post-disagreement state epistemic properties regarding p are lesser on R than B 's. Does epistemic improvement

automatically take place, simply because *A* conciliates? The answer is ‘no’. To see this, consider an easy case. Suppose that, post-disagreement, *A* holds a false belief that not-*p* with some degree of justification (say, 0.6), while *B* more justifiably (0.7) truly believes that *p*. By many rankings *R* compatible with extant epistemological views, *B* is doing better than *A*. At the same time, there is no guarantee that if *A* permissibly conciliates, *A* comes to improve her state epistemic properties with regard to the target proposition by all *R* rankings.

To see this, take a case in which *B* is a three-year-old, Ben, while *A* is a regular adult, Anna, and the proposition at stake concerns the chemical formula of sulphonic acid. In such a situation, EIDN₃, in conjunction with a justification-centric ranking *R*, would recommend a conciliatory response. After all, Ben has a higher degree of justification than Anna. However, it would seem that, were Anna to conciliate in the face of Ben’s testimony, she would, at best, acquire a true belief that is justified to a fairly minimal degree—after all, the weight of child testimony on matters of inorganic chemistry can’t go very far.

What this suggests is that the ranking that we are, in fact, interested in is the ranking of state properties post-disagreement versus state properties post-conciliation *for the same disagreeing party*, rather than the ranking of state properties post-disagreement between disagreeing parties. That is, what matters is not whether, post-disagreement, my doxastic state enjoys a better set of properties than yours, but rather that, in light of the disagreement, the set of epistemic properties my doxastic state enjoys improves from *t*₁ (the time after disagreement occurs) to *t*₂ (the time after conciliation in virtue of the disagreement). If it does, I should conciliate. If it does not, I should hold steadfast.

By way of illustration: in the previous case, what we are interested in, as far as epistemic improvement is concerned, is how the post-disagreement state epistemic properties of Anna’s doxastic attitude rank with respect to the state properties of the doxastic attitude Anna would form were she to conciliate in virtue of her disagreement with Ben. On the justification-centric ranking *R* that we are working with in this case, Anna should hold steadfast, because, were she to conciliate, her doxastic attitude’s state properties would be downgraded rather than improved: unlike trusting the testimony of an expert on inorganic chemistry, trusting the word of a three-year-old doesn’t simply improve the degree of justification of one’s beliefs on those matters, but it rather downgrades it. Since epistemic improvement towards the goal of having justified beliefs is not possible for Anna in her current situation, she does well to stick to her guns. (Conversely, though, on a veritist ranking that primes truth over justification, the recommendation for this case would be the opposite: Anna should trust Ben’s word because in this way she would acquire a true belief, no matter how unjustified it is).

The bottom line, then, is that whatever ranking of sets of state epistemic properties one considers correct, the parties in a disagreement should always—as per EIDN’s clause (i) below—aim for an improvement in state epistemic properties towards the relevant epistemic goal *G* of the operative ranking, whenever possible—an improvement in state properties of their post-conciliation attitudes with respect to those of their post-disagreement attitudes. If not possible, either because conciliation would mean losing the better state properties one’s post-disagreement attitude already enjoys, or else because shifting from believing one proposition to the other would score as low on the operative scale (such as on a purely justificationist ranking), what one

should do—as per EIDN’s clause (ii) below—is simply to hold steadfast, thus avoiding epistemic demotion (in the worst case scenario) or at least no demotion but lack of epistemic improvement by conciliation (in the less detrimental yet suboptimal case). Here is, then, the final formulation of the norm we want to propose given these insights:

The Epistemic Improvement Norm of Disagreement (EIDN): Given a value ranking R of sets of state epistemic properties mapped with respect to proximity to epistemic goal G , for all cases in which A and B disagree about whether p —where, after the discovery of the disagreement, A has a doxastic attitude with content p with a set of epistemic properties E_1 and B has a doxastic attitude with content not- p with a set of epistemic properties E_2 , A *prima facie* should:

- (i) conciliate in virtue of her disagreement with B if and only if A ’s doxastic attitude thus adopted would thereby enjoy a set of epistemic properties E_3 , such that E_3 ranks higher than E_1 on R ,
- (ii) or else hold steadfast in virtue of her disagreement with B .

According to EIDN, then, the thing to do in disagreement cases (as in our epistemic lives in general) is to improve epistemically as much as possible and, if this is not an option, one should at least adopt an epistemic damage control policy by holding steadfast thereby avoiding epistemic demotion or at least costly conciliation when lack of epistemic improvement is not available (see fn. 11).

One final interesting thing to note, before seeing EIDN at work, is that it scores highly on *informativeness*. To be sure, for philosophers with internalist intuitions, in search for perfect luminous conditions, EIDN will likely be just as frustrating as KDN was: just like KDN, EIDN is not perfectly operationalizable. But we have seen that no plausible norms are, and that no plausible non-trivial conditions are perfectly luminous, so we are not bothered by this result. Crucially, though, in contrast to KDN, EIDN is fully informative: it takes an epistemic goal G and a ranking R as input, and it outputs a recommendation as to how to proceed to achieve enhanced closeness to G on R . In this, EIDN does not suffer from sameness in goal and norm. Just like, e.g., the traffic norm ‘Drive at most 50 km/h within city bounds!’ with the respect to the aim of achieving safe driving, EIDN is not offering a perfectly luminous condition for norm compliance, but it is informative with regard to what one is to do to make progress towards one’s epistemic goals. Of course, that being said, KDN itself delivers a particular ranking of sets of epistemic properties, which can also, in turn, be plugged into EIDN, and which would result in identical outcomes for the two norms.

6 EIDN at work

To see how EIDN would work when coupled with the main views on the normativity of belief on the market, we will look at its performance on one central case discussed in the peer disagreement literature, as well as a simple case of everyday disagreement. Here are the cases:

Restaurant Check

Suppose that five of us go out to dinner. It's time to pay the check, so the question we're interested in is how much we each owe. [...] I do the math in my head and become highly confident that our shares are \$43 each. Meanwhile, my friend does the math in her head and becomes highly confident that our shares are \$45 each. (adapted from Christensen, 2007: 193)

Everyday Disagreement

Ann is my climate-change-denier friend: I know climate change is happening, while Ann unjustifiably, against all evidence, believes that it is not.

Let's consider how EIDN would work in conjunction with two popular views on the normativity of belief on the market—one that takes truth to be the chief epistemic good, and another that favors knowledge.

Suppose, first, that you are a traditional, truth-first virtue reliabilist about the normativity of belief, along the lines first developed by Sosa (e.g., 2007). On (a very simplified version of) this view, the normativity of belief is modelled on general performance normativity: an AAA-model of performance assessment.¹³ We can assess performances for accuracy, adroitness, and aptness. Accurate performances achieve their aim, adroit performances manifest competence, and apt performances are accurate in virtue of being adroit: their success manifests competence. On this view, then, apt performance is normatively superior to both adroit and accurate performance. In turn, for the case of belief, apt performance will amount to knowledge (accuracy in virtue of competence), adroit performance will amount to justified belief (competent belief), while truth will be the accuracy condition. One way compatible with AAA (though not the only one) to spell out the relevant ranking of epistemic state properties, then, would be to rank apt belief as better than adroit belief, adroit belief as better than accurate belief, and, finally, accurate (i.e., true) belief as better than false belief.

In contrast, a simple knowledge-centric picture like the one H&S favour will likely predict different epistemic value rankings. We have seen in the discussion above that a plausible knowledge-first normative picture needs to make good on the thought that some ways of being ignorant are better than others, and also that withholding is not more valuable than some of these ways. In this spirit, let's then take a toy ranking organized around closeness to knowledge along the lines of: knowledge > true belief > evidenced credence > suspension > false belief.

Let's now return to the cases. In *Restaurant Check*, on the truth-first virtue reliabilist ranking we have assumed, both myself and my friend are justified to believe the initial results of the calculus we perform at t_1 —i.e., before the disagreement becomes known. After all, we employed a truth-reliable competence to produce these beliefs. If so, I have an adroit (justified) belief that the correct result is \$43 (p), while my friend can easily reason herself from her adroit (justified) belief that the correct result is \$45, to a justified belief that the correct result is not \$43 (not- p). Furthermore, defenders of this view accept epistemic defeat. On the vast majority of views of defeat, after having registered this disagreement as to whether the correct result is \$43 or not, our initial justification is, at least to some extent, defeated. If so, the view will predict, in line with conciliationist intuitions (which this case was put forth to support to begin with), that maintaining our full beliefs, respectively, that p /not- p is impermissible for

¹³ For related but subtly different approaches, see Greco (2003, 2010) and Broncano-Berrocal (2017, 2018).

my friend and I. In a nutshell, we should both reduce our confidence in our respective beliefs. It seems plausible then that, at most, what we are left with post-disagreement will be justified credences that fall short of full belief,¹⁴ or partially justified full beliefs. Now, crucially, the important asymmetry of state epistemic properties kicks in: say that my belief/credence is true, while my friends' is false. What does EIDN predict, in conjunction with virtue reliabilism? Neither of our beliefs are apt after defeat. However, my belief, as opposed to my friends', is accurate. Virtue reliabilism, in conjunction with EIDN, makes two important, extensionally adequate predictions: first, that conciliationists are right in thinking that holding steadfast in this case is impermissible for both myself and my friend, i.e., that we should both lower our confidence after defeat; and, second, that since I'm right and she's wrong, she is the one who should change her mind.

Consider now EIDN conjoined with a simple knowledge-centric picture. We have seen that many defenders of simple knowledge-first pictures (e.g., Lasonen-Aarnio, 2014) do not accept knowledge defeat. According to these philosophers, if I know that p , I remain justified to believe that p in the face of misleading evidence (*ibid.*). On this kind of account as well, then, EIDN will predict that, in the *Restaurant Check*, my friend should conciliate—since my belief is knowledgeable and hers is false. However, in contrast to the previous account we looked at, on this combination of views, we need to appeal to an error theory to explain away the conciliationist intuitions: the view predicts I may remain fully confident that p .

Consider now the case of my everyday disagreement with my climate change denialist friend, Ann. EIDN in conjunction with the virtue reliabilist ranking we have assumed for illustration will predict (correctly) that Ann should conciliate: after all, my belief that climate change is real is apt, while hers is inaccurate (false).

As the case is described, EIDN with our knowledge-first toy ranking would also give a good result in this case: I should hold steadfast—after all, I know—while Ann should conciliate because her unjustified false belief ranks lower than knowledge.

Going back to H&S's KDN briefly: of course, KDN itself delivers, again, a particular ranking of sets of epistemic properties, which can also be plugged into EIDN, and which would result in identical outcomes for the two norms. However, it is interesting to note that the prediction delivered by our toy knowledge-first model in conjunction with EIDN is strongly superior to the one that KDN would output. To see this, consider a variation of the everyday disagreement case in which Ann just cannot get herself to fully believe, as a result of our disagreement, that I am right (which is also the most typical way this scenario would go). In this case, KDN predicts she should withhold after we interact. After all, she cannot get herself to fully believe, and thereby she cannot know as a result of this disagreement that climate change is happening. In contrast, EIDN in conjunction with our toy model predicts that, were she in a position to at least acquire an evidenced true credence that climate change is happening, she should go for it. This is also, intuitively, the correct result in real cases of disagreement: plausibly, what we aim for when disagreeing with someone like Ann, is to move them as much as possible in the direction of the truth, or knowledge.

¹⁴ For simplicity's sake, we ignore complications pertaining to the relation between credences and full beliefs here.

In a nutshell, EIDN's predictions will be just as extensionally adequate as the theory of the normativity of belief it is conjoined with will allow. And that is exactly how it should be—if we want our norm of belief in the case of disagreement to map on neatly to our general account of the norm of belief.

7 Concluding remarks

This paper has proposed a *methodological turn* for the epistemology of disagreement: instead of starting with highly idealized cases of peer disagreement, we have focused on disagreement *simpliciter*. In turn, we have proposed a norm of disagreement on which, very roughly, what one should do when faced with a disagreeing party is to improve the epistemic properties of one's doxastic attitude or, alternatively, hold steadfast. The norm we have proposed can accommodate any ranking of state epistemic properties, sourced in any view on the normativity of belief on the market, from hard-core internalist positions to extreme externalisms. In this, our proposed norm opens up a new field of research into the question of what ranking of state epistemic properties is more extensionally adequate. The answer to this question will tell us which view of the normativity of belief handles cases of disagreement better.

Acknowledgements Many thanks to Chris Kelp and two anonymous referees for extensive feedback on this paper. Thanks also to the audiences at the Higher Seminar in Theoretical Philosophy (Lund University), the Seminar of the Research Unit for Epistemology, Metaphysics, and Philosophy of Cognition (Aarhus University), The Social Epistemology Research Group Seminar (University of Copenhagen), the Social Epistemology Workshop (University of Bristol), and the Complex Disagreement Workshop (University of St Andrews). This research has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 948356) and from a 2019 Leonardo Grant for Researchers and Cultural Creators awarded by the BBVA Foundation.

Funding Open Access funding provided thanks to the CRUE-CSIC agreement with Springer Nature.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Bergmann, M. (2009). Rational disagreement after full disclosure. *Episteme*, 6, 336–353.
- Bogardus, T. (2009). A vindication of the equal-weight view. *Episteme*, 6, 324–335.
- Broncano-Berrocal, F. (2017). A robust enough virtue epistemology. *Synthese*, 194(6), 2147–2174.
- Broncano-Berrocal, F. (2018). Purifying impure virtue epistemology. *Philosophical Studies*, 175(2), 385–410.
- Chisholm, R. M. (1963). Contrary-to-duty imperatives and deontic logic. *Analysis*, 24(2), 33–36.

- Christensen, D. (2007). Epistemology of disagreement: The good news. *Philosophical Review*, *116*, 187–217.
- Christensen, D. (2009). Disagreement as evidence: The epistemology of controversy. *Philosophy Compass*, *4*, 756–767.
- Christensen, D. (2010). Higher order evidence. *Philosophy and Phenomenological Research*, *81*, 185–215.
- Christensen, D. (2013). Epistemic modesty defended. In D. Christensen & J. Lackey (Eds.), *The epistemology of disagreement: New essays*. Oxford: Oxford University Press.
- Cohen, S. (2013). A defense of the (almost) equal weight view. In D. Christensen & J. Lackey (Eds.), *The epistemology of disagreement: New essays*. Oxford: Oxford University Press.
- Decker, J. (2014). Conciliation and self-incrimination. *Erkenntnis*, *79*, 1099–1134.
- Elga, A. (2007). Reflection and disagreement. *Noûs*, *4*, 478–502.
- Elga, A. (2010). How to disagree about how to disagree. In T. Warfield & R. Feldman (Eds.), *Disagreement*. Oxford: Oxford University Press.
- Feldman, R. (2005). Respecting the evidence. *Philosophical Perspectives*, *19*, 95–119.
- Feldman, R. (2009). Evidentialism, higher-order evidence, and disagreement. *Episteme*, *6*, 294–312.
- Feldman, R. (2007). Reasonable religious disagreements. In L. Antony (Ed.), *Philosophers without gods: Meditations on atheism and the secular* (pp. 194–214). Oxford: Oxford University Press.
- Frances, B. (2010). The reflective epistemic renegade. *Philosophy and Phenomenological Research*, *81*, 419–463.
- Frances, B. (2014). *Disagreement*. Cambridge: Polity.
- Geach, P. (1956). Good and evil. *Analysis*, *17*, 33–42.
- Hawthorne, J., & Srinivasan, A. (2013). Disagreement without transparency: Some bleak thoughts. In D. Christensen & J. Lackey (Eds.), *The epistemology of disagreement: New essays*. Oxford: Oxford University Press.
- van Inwagen, P. (2010). We're right. They're wrong. In T. Warfield & R. Feldman (Eds.), *Disagreement*. Oxford: Oxford University Press.
- Kelly, T. (2010). Peer disagreement and higher-order evidence. In R. Feldman & T. A. Warfield (Eds.), *Disagreement* (pp. 111–174). Oxford: Oxford University Press.
- Kelly, T. (2005). The epistemic significance of disagreement. In T. Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology* (Vol. 1, pp. 167–196). Oxford University Press.
- Kornblith, H. (2010). Belief in the face of controversy. In R. Feldman & T. A. Warfield (Eds.), *Disagreement*. Oxford: Oxford University Press.
- Kornblith, H. (2013). Is philosophical knowledge possible? In D. E. Machuca (Ed.), *Disagreement and skepticism*. London: Routledge.
- Lackey, J. (2008). What should we do when we disagree? In T. S. Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology* (pp. 274–293). Oxford: Oxford University Press.
- Lackey, J. (2010). A justificationist view of disagreement's epistemic significance. In A. Haddock, A. Millar, & D. Pritchard (Eds.), *Social epistemology*. Oxford: Oxford University Press.
- Lackey, J. (2013). Disagreement and belief dependence: Why numbers matter. In D. Christensen & J. Lackey (eds.), *The epistemology of disagreement: New essays* (pp. 243–268). Oxford, UK: Oxford University Press.
- Lasonen-Aarnio, M. (2014). Higher-order evidence and the limits of defeat. *Philosophy and Phenomenological Research*, *88*, 314–345.
- Matheson, J. (2015). *The epistemology of disagreement*. Basingstoke: Palgrave.
- McHugh, C. (2012). The truth norm of belief. *Pacific Philosophical Quarterly*, *93*, 8–30.
- McKinnon, Rachel. (2015). *The norms of assertion truth, lies and warrant*. New York: Palgrave Macmillan.
- Pittard, J. (2015). Resolute conciliationism. *Philosophical Quarterly*, *65*, 442–463.
- Rabinowicz, W., & Rønnow-Rasmussen, T. (2004). The strike of the demon: On fitting pro-attitudes and value. *Ethics*, *114*, 391–423.
- Simion, M. (2021). *Shifty speech and independent thought*. Oxford: Oxford University Press.
- Simion, M., Kelp, C., & Ghijssen, H. (2016). Norms of belief. *Philosophical Issues*, *26*(1), 375–392.
- Titelbaum, M. (2015). Rationality's fixed point (Or: In defense of right reason). In T. Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology* (Vol. 5, pp. 253–294). Oxford: Oxford University Press.
- Vavova, K. (2014). Confidence, evidence, and disagreement. *Erkenntnis*, *79*, 173–183.
- Weatherston, B. (2013). Disagreements, philosophical and otherwise. In D. Christensen & J. Lackey (Eds.), *The epistemology of disagreement: New essays*. Oxford: Oxford University Press.

- Wedgwood, R. (2010). The moral evil demons. In T. Warfield & R. Feldman (Eds.), *Disagreement*. Oxford: Oxford University Press.
- Weintraub, R. (2013). Can steadfast peer disagreement be rational? *Philosophical Quarterly*, 63, 740–759.
- White, R. (2005). Epistemic permissiveness. *Philosophical Perspectives*, 19, 445–459.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.