

Knowledge and Disagreement

Draft

Mona Simion (Cogito Epistemology Research Centre, University of Glasgow) & Fernando Broncano-Berrocal (LOGOS Research Group, University of Barcelona)

1 Introduction

Disagreement occupies a significant place in our lives, from mundane disagreements (e.g., is pineapple pizza tasty?) to very deep, entrenched ones (e.g., is death penalty permissible?). The significance of disagreement is practical, political, moral, psychological but, also, epistemic. When we disagree with each other, we might not only be interested in knowing the impact such conflicts have on, e.g., the legitimacy of our political decisions, the possibility of assessing some moral claims for truth, or our psychological well-being, but also on how we should go about forming our beliefs.

Accordingly, the key question that the literature on the epistemology of disagreement promises to answer is the question of what is permissible for one to believe in the face of disagreement (Lackey 2010a). However, this promise, far from having materialized, has simply become less ambitious. In particular, the debate has almost only focused on a particular kind of disagreement: *peer disagreement*, which is the somewhat idealized disagreement between so-called epistemic peers, i.e., agents who possess roughly the same evidence and cognitive capacities. Thus, the less ambitious promise has been to provide a satisfactory answer to this other question: What would it be rational for two acknowledged epistemic peers to believe upon discovering that they are in disagreement? (Frances & Matheson 2019).

Unsurprisingly, most views of the epistemology of disagreement have focused on analyzing idealized rather than real-life cases of disagreement with the hope that this can help reduce the noise of the latter. The ambition, then, is that we can subsequently upload context – i.e., add the details of particular cases to our theoretical results – and generalize our results to cases of real-life disagreement. In this way, the literature has been driven by a methodological approach to disagreement that puts the analysis of peer disagreement methodologically first (henceforth *Peer-Disagreement-First*, or PDF) in the analysis of the normativity of belief in the face of disagreement.

Without any doubt, PDF has been fruitful on several scores. For example, thanks to PDF, we now have a better grip on significant epistemological questions, including the role that higher-order evidence

(Kelly 2010; Lasonen-Aarnio 2014), personal information (van Inwagen 2010) or justification (Lackey 2010b) play in disagreement cases, or the question of whether evidence is permissive, i.e., whether it can rationalize opposite doxastic attitudes (Feldman 2007). Crucially, though, the more general, more ambitious question of what is permissible for one to believe in the face of disagreement remains unanswered.

This paper does three things: first (§2), it argues that the prospects of ‘stripping off’ the idealizations, and, ultimately, of giving an answer to the general question are dim, so long as the only methodological approach to the epistemology of disagreement remains PDF.

The second aim of the paper is to investigate the prospects of an alternative, knowledge-first approach to disagreement. This approach takes knowledge to be the central value of the epistemic domain, and norms governing moves in this domain – such as belief in the face of disagreement – to drop right out of this value. We first (§3) look at a knowledge-first view of disagreement defended by John Hawthorne and Amia Srinivasan (2013), and argue that it remains unsatisfactory at two normative junctures. Second, we put forth our view of the normative structure of disagreement cases – according to which, roughly, when faced with disagreement as to whether something is the case, one should improve one’s doxastic status as much as possible, or else hold steadfast. Last, we show how a knowledge-first variety of this view will be strongly superior to the competition (§4). Finally, in §5 we conclude.

2 Peer-Disagreement-First

2.1 Two Assumptions Behind PDF

To begin with, it will be useful to remind ourselves of the motivations behind the peer-disagreement-first approach. Recall that the thought was that, by only focusing on idealized cases, we can reduce the noise of real-life disagreement, and thus have better access to the data, and increase the likelihood of extensional adequacy. The idea, then, is that we can subsequently upload context – i.e. add the details of particular cases to our theoretical results – and thereby generalise our results to cases of real-life disagreement. Here is David Christensen:

The hope is that by studying this sort of artificially simple socio-epistemic interaction, we will test general principles that could be extended to more complicated and realistic situations, such as the ones encountered by all of us who have views—perhaps strongly held ones—in areas where smart, honest, well-informed opinion is deeply divided. (Christensen 2009: 231).

Even more optimistically, Jonathan Matheson claims:

Idealized cases of disagreement can set the ‘baseline’ for the epistemic significance of disagreement in general. How we

should respond to evidence about everyday, real-world disagreement will depend at least in part on how we should respond to idealized disagreement (even if we never in fact encounter idealized disagreements). Once such a baseline response is set, the epistemic significance of disagreement in everyday cases can be discovered by examining in what ways the everyday case differs from the idealized case. (Matheson 2015: 33-34).

In sum, then, the two main assumptions that justify the PDF approach are:

Easy Extensional Adequacy Claim: By focusing on simple, idealized, peer-disagreement cases, we have an increased likelihood of extensional adequacy.

Easy Context Upload Claim: Generalizing from permissibility facts for idealized cases to permissibility facts for cases of real-disagreement via context upload is easier than directly identifying permissibility facts for cases of real disagreement.

In what follows, we take issue with both of these claims. All in all, we take it, we should consider methodological revision.

2.2 The Problem of the Instability of Intuitions

Recent work in experimental philosophy (Alexander et al. (2018)) provides empirical evidence that the PDF methodology faces insurmountable difficulties when it comes to evaluating extensional adequacy via the method of cases. What Alexander and colleagues' survey shows is that people have conciliatory or steadfast intuitions about the cases depending on what questions they are asked—this is known as a *framing effect*. In particular, when they are asked to make comparative evaluations (e.g. 'should you always give your peer's opinion equal weight?') subjects have more steadfast responses than when they are asked to make non-comparative judgments (e.g. how confident should you now be in your opinion after realizing that you are in a disagreement with a peer?). This *instability of intuitions problem* suggests, contra the *Easy Extensional Adequacy Claim*, that PDF does not offer increased likelihood of extensional adequacy.

The following two problems shed doubt on the *Easy Context Upload Claim* underlying PDF.

2.3 The transition problem

The transition problem is the problem of generalizing one's view of peer disagreement to real-life disagreement. We have two main reasons to be skeptical about the successful normative transition from peer disagreement cases to real-life cases: the first pertains to the way PDF is construed in the literature; the second is a more general, theoretical worry.

2.3.1. *The context upload worry*

Extant PDF theories are mostly silent about how to properly transition to real-life cases of disagreement. Here is why: say that you are a conciliatorist. According to you, roughly, in the face of peer disagreement, you should (in one way or another) weaken your confidence in the relevant proposition. Now let us try to generalize to two real-life cases. The first one is a case of expert-to-non-expert disagreement. I (non-expert) believe that p , while you (expert on the topic at hand) believe that not- p . Transition in such cases does not constitute any difficulty for the view: conciliationism affords a straightforward answer here. If peer testimony is enough for a conciliatory response to be the right one, surely expert testimony should do the trick too. The problem, however, arises when the cases under discussion are cases in which I am disagreeing with someone who is less than my peer: in particular, *slightly* less than my peer. Such cases are extremely ubiquitous in everyday life: we tend to talk a lot to people we share a lot with. Say that my friend, Ann, disagrees with me on the exact percentage of people vaccinated against Covid-19 at a specific time. She's only slightly less well informed than me on the issue – say she has less time than I do to read the media. What should I believe in the face of disagreement? Conciliationism offers no clear-cut *principle* stating what is permissible or impermissible to do in this kind of cases.

Obviously, conversely, steadfast views have no trouble giving a verdict in Ann's case: since, according to them, I should hold on to my beliefs in cases of peer disagreement, surely I am entitled to do the same when I disagree with someone who's less than my peer.

However, it is unclear what, on such views, the *principled* rational response to cases of expert disagreement should be, or even to everyday cases in which I disagree with someone who is slightly more informed than myself.

2.3.2 *The 'Car in the Dark' Worry*

One can be misled to view PDF as quite a scientific methodological approach: by starting with the idealized peer disagreement scenario, it isolates itself from the real-life noise by taking the epistemology of disagreement from the outside world to the lab.

One important thing to note about this methodology, though: when scientists get rid of the contingent contextual features in order to test the behavior of the relevant target phenomenon in isolation, they remove noise but make sure to preserve the *phenomenon itself* intact. Crucially, on the vast majority of extant models of scientific idealization,¹ this implies some restrictions on what properties are permissibly idealized from.

It's not clear, though, whether this methodology is properly employed by PDF: after all, since we are dealing with a normative phenomenon, an analogous approach to the scientific one would be to

¹ For an excellent taxonomy see Weisberg (2009).

get rid of normative contextual noise, but preserve the normative core of the phenomenon. Otherwise, there is not much left to be tested.

One way in which epistemologists and normative philosophers in general usually do this is by stipulating their ambition merely to provide *prima facie*, *pro tanto* normative accounts. This isolates from defeat, normative overriding at the context, and so on. PDF does more than that, though: by stipulating that the parties in the disagreement are perfect (or nearly perfect) epistemic peers, it threatens to strip the disagreement scenario of far more normative thickness than that, if not even of most/all normative constraints characteristic to the phenomenon of disagreement itself.

To see this worry, consider the following three scenarios: you are in a town you've never been in before, about which you don't know much. You are asked to observe traffic and figure out what traffic norms hold in the town in question. In scenario 1, you are asked to do this while driving through town during heavy traffic, while a hurricane and a terrorist attack are ongoing. In scenario 2, you are asked to watch a video of a simulation of traffic in the town in question. In scenario 3, you are driving your car in complete darkness in the same town. We submit that 2 improves over 1 in that it gets rid of the noise. We also claim that 3 is the worst scenario. The difference between scenario 2 and 3, of course, is that in 2, but not in 3, the normative phenomenon to be investigated is still accessible to the investigator in its fullness. The worry here is that PDF goes too far in the direction of 3; *prima facie* and *pro tanto* provisos are enough for isolating from contextual noise. Stripping disagreement cases of all their epistemic properties is, we think, taking the methodology one step too far in the direction of trying to theorize in the dark.²

2.4 The normative fit problem

It is worth noting once more that the central question in the disagreement literature – ‘What is permissible for one to believe in the face of disagreement?’ – is but an instance of the more general question, central to epistemology at large, ‘What is permissible to believe?’. The literature concerned with the latter does not employ much in the way of idealizations, that is, apart from the fact that what is usually discussed is *prima facie* permissibility: several nicely worked out accounts thereof are available on the market, and the debate is ongoing.³ Since this more general literature on the normativity of belief does not look at idealized scenarios, it does not need to transition from one set of cases to another, as current views in the literature on disagreement are compelled to do. What the more general literature does is to propose *norms of belief* applicable to the whole spectrum of cases of believing.⁴

This being so, why is it that, in the case of disagreement alone, we need a sterile epistemic environment for being able to answer the question of the normativity of belief, while the entire literature involved

² We owe the analogy to Chris Kelp (pc).

³ See Chignell 2018 for an overview.

⁴ See, e.g., Simion et al. 2016.

in giving an answer to the more general question does not seem to be bothered by the noise of ordinary cases of belief? Second, and more importantly, how are we to reconcile the two approaches? That is, how will the results of the two lines of research *fit together*? This is what we call the *normative fit problem*.

A possible reply to this worry would insist that believing in disagreement situations is somehow *special*, i.e., governed by norms pertaining to this particular setting. After all, just because belief in disagreement situations is but an instance of belief in general, it need not be governed by one and the same norm. In the same way, just because waltzing is a species of dancing, it need not follow that it is merely governed by norms pertaining to dancing in general.

However, this line of thinking cannot be correct. Normative inheritance holds from type to species, in that for all S, if S is a species of type T, then S is governed by norms that are at least as strong as those governing T (Simion 2018). Short of inheriting T's norms, S is not a species of T. True, what makes waltzing a particular species of dancing to begin with is the fact that it is governed by extra-norms, *on top of the ones governing general dancing*. But if it is not governed by the norms that govern dancing, waltzing is not a species of dancing to begin with. Similarly, the norm for belief in disagreement situations may be stronger than the norm for belief in general, or it might be that more norms govern it. To see this, say, for instance, that you advocate a knowledge norm of belief, *à la* Williamson (2000)—one's belief that *p* is epistemically permissible only if one knows that *p*. Compatibly with that, you might think that, even if you know it is raining outside, testimony to the contrary from an epistemic peer is good enough reason to check again. One reason why this might be so is because knowledge gets defeated by your peer's testimony.

If that is the case, though, it is rather surprising that no more has been said about how the two fields of research interact, and this is plausibly due to the fact that, from day one, the literature has adopted a methodology that makes the analysis of one set of cases (peer disagreement) a distinct affair from the analysis of another set (real-life disagreement).

3 Epistemology of Disagreement, Knowledge-First

In what follows, we want to argue that the (social) epistemology of disagreement should take a cue from the literature on individualist epistemology, in which a novel methodology has been gaining significant traction recently: a knowledge-first approach to epistemological affairs. The most prominent champion of this kind of view is Timothy Williamson (2000), who has pioneered the approach to epistemology. One key thought here is that rather than trying to analyze knowledge in terms of various other epistemic phenomena, such as justification, evidence or understanding, other epistemic phenomena are to be analyzed in terms of knowledge.

More specifically, in what follows, we want to argue that one important social epistemological aspect of disagreement is omitted by the

mainstream, PDF approach: disagreement’s epistemic function. Why is it that we bother to engage in disagreements with each other to begin with? To what effect? Here is one plausible hypothesis the credentials of which we will explore henceforth: the epistemic function of disagreements is to get us closer to knowledge. The practice of disagreeing with each other, on this view, while it may well be rather unpleasant at times – and thus prudentially to be avoided – is often rendered all-things-considered worth it due to its epistemic advantages: in disagreeing with one another, we come closer to the goal of our inquiries. If this is so, we can do *knowledge-first* epistemology of disagreement: we can analyze the permissible doxastic response in the face of disagreement in terms of knowledge. This methodology, in turn, would have at least two major advantages over the PDF competition: since the view would regard general disagreement as the focal point of evaluation, it would face no transition problem; since the knowledge-first literature features well-developed accounts of the normativity of belief, evidence, and justification, the view would enjoy exceptional normative fit – and thereby would inherit prior plausibility from the general framework.

Before putting forth our own preferred knowledge-first recipe for disagreement, in what follows, we take a brief look at a promising predecessor: John Hawthorne’s and Amia Srinivassan’s *Knowledge Norm of Disagreement* (2013).

3.1 The Knowledge Disagreement Norm

Hawthorne and Srinivasan (2013) (henceforth H&S) have put forward a proposal according to which the appropriate response to disagreement, as with all other epistemic activities we engage in, is to aim to gain *knowledge*. On their view, “[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). Given this, the core of their account is a norm of disagreement according to which the rational thing to do in the face of disagreement is the following:

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*.

KDN is claimed to follow from an approach to epistemology on which the most fundamental epistemic concern is knowledge. In particular, the following claim is widely subscribed by advocates of such an approach:

The Knowledge Norm of Belief (KNB): One's belief that p is epistemically permissible only if one knows that p .

The upshot, then, is that the norm of belief in disagreement cases, KDN, becomes an instance of the general norm of belief, KNB.

We are generally sympathetic to the picture painted by H&S for the epistemology of disagreement. We have argued extensively in previous work (Broncano-Berrocal & Simion Forthcoming), however, that the particular norm of disagreement that they propose is problematic in more than one way. We will not rehearse all these arguments here, but rather limit ourselves to discussing the worries that KDN faces from within the knowledge-first framework.

We have two main worries for KDN: the first is a purely theoretical one: we worry that, contra H&S, it is not the case that KDN 'falls out' of a knowledge-centric way of doing epistemology. As we are about to argue, to the contrary, there are several ways in which it sits quite uncomfortably within such an approach. Second, we worry that KDN makes implausible predictions for the vast majority of real-life disagreement cases due to the phenomenon of defeat. In what follows, we discuss these two worries in turn.

3.1.1 The Prior Plausibility Worry

We worry that, if knowledge is, indeed, the main telos of all of our epistemic endeavours, KDN paints a very strange normative-theoretic picture, whereby the norm and goal of belief coincide. This doesn't happen a lot in the normative landscape, or so we argue. To see this, it will be useful to remind ourselves what motivated KDN to begin with. First, H&S endorse:

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

Furthermore, H&S also assume that, given that in a knowledge-centric epistemology knowledge is taken to be the main epistemic goal, we should only believe what we know, and otherwise withhold. Here it is:

The Knowledge Goal—Knowledge Norm Link (KGKN Link): If KGT is true, one should only believe what one knows.

Note that KGKN Link is, *de facto*, a conditional prior plausibility claim: should one endorse a knowledge-first epistemological picture, one should be sympathetic to KDN. We are suspicious of KGKN Link. To see why, think of other norms in other normative domains. Take, for instance, norms governing traffic. Plausibly, their aim is to generate safe traffic. Now, note that, to this effect, these norms do not share content with the goal in question: they don't just say 'Drive safely!' to the aim of driving safely. That would be rather uninformative. Instead, they prescribe means that, in normal conditions, are reliably conducive to reaching the aim of safe traffic: 'Drive 30 miles/hour!' is one such traffic

norm. Other norms all over the normative landscape follow suit. Now, the point we are trying to make is that, on grounds of prior plausibility, we should expect the epistemic normative domain to exhibit similar behaviour, in that norms prescribe informatively reliable means of reaching the goal(s) of the domain. For this to be the case, goals and norms should not share content. Epistemic norms governing belief formation should prescribe reliable ways to form beliefs such that the aim of belief – knowledge – is reliably met. A norm that tells us to only believe what we know is not very informative when it comes to forming knowledgeable beliefs.

Let's return to KDN. If one is a knowledge-centric epistemologist – i.e. someone that thinks that knowledge is the aim of believing – then they should resist a knowledge norm of belief, since it turns the epistemic into a normative oddball in the normative landscape: on pain of loss in prior plausibility, a knowledge aim of belief does not sit well with a knowledge norm of belief. As such, we have reason to be suspicious of KGKN Link. But if that is the case, KDN loses its main rationale, in that it does not straightforwardly drop out of a knowledge-first picture.

3.1.2 *The Defeat Worry*

Many times, if I justifiably believe that p , and you come and tell me that not- p , my initial justification for believing p will be defeated (unless I have reason to believe you don't know what you're talking about). Indeed, something in the vicinity of this claim is the direct result of pretty much any view ever defended in the epistemology of testimony: testimony from a source that you have reason to trust provides (*pro tanto*, *prima facie*) justification. If I have justification to believe that not- p , though, sourced in your testimony that not- p , surely my justification to believe that p will be correspondingly affected: after all, that's how the balance of reasons plausibly goes in normativity in general. I should, then – at least – lower my confidence in p , should you tell me that not- p . Of course, if this is so, the possibility of acquiring knowledge will not be present in the vast majority of disagreement cases. If so, it would seem that KDN would recommend generalized withholding as the one and only permissible response to the vast majority of disagreements. This seems extreme; also, it gives the knowledge norm a sense of vacuity, since there are no cases – or, at best, very few – in which one can end up knowledgeable after disagreement, and thus knowledge has very little to do with what KDN ends up prescribing – i.e. generalized withholding.

However, defenders of knowledge-first approaches tend to reject knowledge defeat (see, e.g., Lasonen-Aarnio 2014) and this includes H&S. These epistemologists' main claim is that holding on to one's beliefs in light of evidence to the contrary is something one is entitled to, provided that one knows. The way they meet the intuitive concerns of impermissible dogmatism that might arise is by proposing error theories of the following form: if we think that knowers that dismiss countervailing evidence are being impermissibly dogmatic when sticking

to their guns, it is because the dismissal of such evidence is epistemically blameworthy, but not because it is epistemically impermissible.

We are sceptical about this move, for two main reasons: in terms of prior plausibility, this move reduces the scope thereof quite drastically. Knowledge-first epistemology is a very successful research programme. Defeat-denying knowledge-first epistemology, however, much less so, for many reasons which we don't have space to fully rehearse here (but see Brown 2018 for an excellent discussion). To mention just three very serious reasons to doubt knowledge defeat skepticism is correct: (1) the claim generalizes implausibly from knowledge defeat to all justification defeat, on a traditional JB=K Williamsonian picture; (2) at a minimum, one needs to countenance some normative load for evidence against p in one's epistemology. As soon as one does so, it is hard to resist normative load for defeat; (3) defeat skepticism comes in conflict with all epistemologies of testimony we currently have, in that it implies knowers can't get justification from the testimony of others. We find all these reasons to reject defeat skepticism, as well as others put forth in the literature, hard to resist. *Mutatis mutandis*, we think KDN is not giving us a very interesting prediction for the vast majority of disagreement cases, which are, or so we claim, cases in which (*pro tanto*) defeat is instantiated.

To sum up: we think a knowledge-first approach to the normativity of belief in the face of disagreement is on the right track. We also think, however, that the simple, knowledge norm of disagreement put forth by Hawthorne & Srinivassan suffers from both theoretical and extensional adequacy-related drawbacks.

In the next section, we will put forth a knowledge-first alternative norm of belief in the face of disagreement. The norm we propose will be grounded in two claims that we find very plausible: that knowledge is, indeed, the telos of our epistemic endeavours, including our practice of disagreeing with each other, and that disagreement's epistemic function is to improve the epistemic status of our doxastic attitudes towards this telos.

4 Disagreement and Closeness to Knowledge

In previous work (Broncano-Berrocal & Simion Forthcoming), we have argued that, when faced with disagreement, subjects should improve the epistemic status of their doxastic attitudes, relative to the main epistemic telos. We find it plausible that inquiry is a telic epistemic practice, and disagreement is a move in inquiry. It is generally the case that moves in practices aim to fulfill the goal of the practice, by making progress towards reaching it. Starting the engine, for instance, is a move in the practice of driving. It generally aims directly at the aim of the practice – i.e. to get you safely and at your destination. Sometimes, moves in practices can also indirectly aim at the goal of the practice: pressing the brake, for instance, aims to slow the car down, which, in turn, is aimed at getting you safely to your destination. Moves in practices aim (directly or indirectly) at reaching the main telos of the practice.

Plausibly, in virtue of being a move in the practice of inquiry, disagreement will be aimed at making progress towards the aim of the

practice. If so, when faced with disagreement, we should make progress towards the epistemic goal of inquiry via epistemically improving our doxastic attitudes. We should aim to adopt the doxastic attitude with the best epistemic properties available, relative to the goal of inquiry. Here is the account in full:

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 .

(Broncano-Berrocal & Simion Forthcoming)

According to EIDN, the thing to do in disagreement cases (as in our epistemic lives in general) is to *make as much progress as possible* towards our epistemic telos and, if this is not an option, at least *avoid epistemic demotion*.

EIDN models the epistemically permissible response to disagreement in terms of the distribution of state epistemic properties (i.e., the epistemic properties of the disagreeing parties' doxastic attitudes), but state epistemic properties can be ranked in different ways. The ambition is for the account to be fully general, namely to apply to all cases of disagreement, and to be compatible with whatever way to rank state epistemic properties one prefers, and with one's choice of epistemic goal. Thus, depending on one's epistemic normative commitments, different answers might follow from EIDN concerning what is permissible to believe in the face of disagreement. We wanted our epistemic improvement account to be compatible with any such commitments.

How does this account fit in a knowledge-first epistemological picture, which takes knowledge to be the central epistemic value, and the goal of our practice of inquiry? Plausibly, there will be more than one way to marry the two accounts, since there will be more than one way to conceive of the relevant value ranking in relation to closeness to knowledge.

Simion (2021) proposes a *Closeness to Knowledge Norm of Disagreement* (CKND), where closeness to knowledge is unpacked in terms of evidential probabilities. What we get for R is:

R_k : S 's doxastic attitude D with content p is closer to knowledge than S^* 's doxastic attitude D^* with content q iff D better matches the evidential probability of p for S than D^* matches the evidential probability of q for S^* .

And the corresponding norm put forth is:

The Closeness to Knowledge Norm of Disagreement (CKND):

In a case in which *A* and *B* disagree about whether *p* —where, after the discovery of the disagreement, *A* has a doxastic attitude *D* with content *p* and *B* has a doxastic attitude *D** with content not-*p*, *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 be closer to knowledge than *D*.

(ii) or else hold steadfast in virtue of her disagreement with *B*.
(Simion 2021)

CKND asks the disagreeing parties to increase their closeness to knowledge (unpacked as an increase in evidential probability) or else stick to their guns. Evidential probabilities are probabilities conditional on one's evidence. The view does not take a stance on the nature of evidence. Just for illustration: A champion on an E=K type of view (Williamson 2000), whereby one's evidence is one's knowledge, for instance, will take *S*'s evidential probabilities to be probabilities conditional on what *S* knows. As such, CKND will predict that what a subject *A* should do when faced with disagreement as to whether *p* is to adopt a doxastic attitude towards *p* that better matches the probability of *p* conditional on what *A* knows, or else hold steadfast. Alternatively, one may prefer a less stringent view of evidence. In previous work, one of us (Simion *Forthcoming*) has defended a milder knowledge-first view of evidence, whereby evidence consists in what one is in a position to know. On this account, CKND will predict that what a subject *A* should do when faced with disagreement as to whether *p* is to adopt a doxastic attitude towards *p* that better matches the probability of *p* conditional on what *A* is in a position to know, or else hold steadfast.

The view has several nice theoretical traits: first, in virtue of being a knowledge-first view, it shares its prior plausibility with the general knowledge-first epistemological framework, and enjoys great normative fit within it. Second, it straightforwardly applies to everyday, non-idealized disagreements, and does so without biting the bullet on the epistemology of defeat. Third, it is informative, in that the norm and goal of disagreement come apart: when aiming at knowing, and faced with disagreement, one should match one's doxastic attitude as closely to one's evidential probabilities as possible. Last but not least, the account is extensionally adequate. To see this, consider two toy cases of disagreement, imported from the PDF literature:

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.)

Maths. I believe $2+2=4$. My neighbour disagrees: according to him, the correct result is 12. (adapted from Lackey 2010a)

The first case is a classic case adduced in support of conciliationism: after all, the thought goes, it seems intuitively unwarranted to stick to my guns in the face of my disagreement with my friend over the restaurant bill. CKND predicts precisely this: after all, my friend's testimony defeats my initial belief. In virtue of defeat, lowering my confidence in my result will better match my evidential probabilities (i.e. probabilities conditional on what I know/am in a position to know etc.) than sticking to my guns.

Maths, conversely, is traditionally used to defend steadfast views of disagreement: it would be strange, the thought goes, to lower my confidence in straightforward analytic truths because of my neighbour's testimony. CKND gets it right again: my neighbor's testimony, although offering *pro tanto* reason to doubt the mathematical truth in question, it does not provide me with strong enough reason to defeat my a priori knowledge: a priori knowledge trumps testimony, on most occasions (although not always – see e.g. Pickel & Schultz; consider also expert testimony about more complicated mathematical truths). My holding steadfast will thus match my evidential probabilities better than conciliating: e.g., given what I know (on undefeated *a priori* grounds), it's more likely that $2+2=4$ than not.

5 Conclusion

We should start talking about disagreement *simpliciter*, rather than about highly idealized cases of peer disagreement, if we want to make further progress in this field. In the spirit of this proposed methodological u-turn, we have proposed that what one should do when faced with (real-life!) disagreement is to seek to improve one's doxastic attitudes with regard to closeness to knowledge.

References

- Broncano-Berrocal, F. and Simion, M. (Forthcoming). Disagreement and Epistemic Improvement. *Synthese*.
- Chignell, A. (2018). The Ethics of Belief. *The Stanford Encyclopedia of Philosophy* (Spring 2018 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/spr2018/entries/ethics-belief/>>.
- Christensen, David (2009). Disagreement as evidence: The epistemology of controversy. *Philosophy Compass* 4: 756-67.
- Feldman, R. (2007). Reasonable religious disagreements. In L. Antony (ed.), *Philosophers without gods: meditations on atheism and the secular*. Oxford University Press: 194-214.
- Frances, B. and Matheson, J. (2019). Disagreement, *The Stanford Encyclopedia of Philosophy* (Winter 2019 Edition), Edward N. Zalta (ed.), URL =

<<https://plato.stanford.edu/archives/win2019/entries/disagreement/>>.

- Hawthorne, J. & Srinivasan, A. (2013). Disagreement without transparency: Some bleak thoughts. In D. Christensen and J. Lackey (eds.), *The Epistemology of Disagreement: New Essays*. Oxford University Press.
- Kelly, T. (2010). Peer disagreement and higher-order evidence. In R. Feldman & T. A. Warfield (eds.), *Disagreement*. Oxford University Press : 111-74.
- Lackey, J. (2010a). What should we do when we disagree? In T. S. Gendler & J. Hawthorne (eds.), *Oxford studies in epistemology*. Oxford University Press: 274-93.
- Lackey, J. (2010b). A justificationist view of disagreement's epistemic significance. In A. Haddock, A. Millar, & D. Pritchard (eds.), *Social epistemology*. Oxford University Press.
- Lasonen-Aarnio, M. (2014). 2014. Higher-order evidence and the limits of defeat, *Philosophy and Phenomenological Research* 88: 314-345.
- Matheson, J. (2015). *The epistemology of disagreement*. Palgrave.
- Pickel, B., and Schultz, M. (2018). Quinean Updates: In Defence of 'Two Dogmas.' *Journal of Philosophy* 115(2): 57-91.
- Simion, M., Kelp, C., and Ghijsen, H. (2016). Norms of belief, *Philosophical Issues*, 26: 375–392.
- Simion, M. (2021). *Knowledge-First Social Epistemology*. Book Manuscript.
- van Inwagen, P. (2010). We're right. they're wrong. In T. Warfield & R. Feldman (eds.), *Disagreement*. Oxford University Press.
- Weisberg, M. (2009). Three kinds of idealization. *Journal of Philosophy* 104: 639-59.
- Williamson, T. (2000). *Knowledge and Its Limits*. Oxford University Press.