Peer Disagreement:
A Call for the Revision of Prior Probabilities

Penultimate draft. Please cite the published version (Dialectica 2015).

The current debate about peer disagreement has so far mainly focused on the question of whether peer disagreements provide genuine counterevidence to which we should respond by revising our credences. By contrast, comparatively little attention has been devoted to the question by which process, if any, such revision should be brought about. The standard assumption is that we update our credences by conditionalizing on the evidence that peer disagreements provide. In this paper, we argue that non-dogmatist views have good reasons to reject conditionalization. Instead, peer disagreements should be understood to call for a revision of our prior conditional probabilities: rather than merely adding to our original evidence, they pose a challenge to the thought that we have properly assessed the probative force of our original evidence.

Keywords: peer disagreement, conditionalization, prior probability, belief revision, certainty, equal weight view

Introduction

The recent debate about peer disagreement has mainly focused on the question of whether, and to what extent, disagreements with one's epistemic peers count as genuine counterevidence to which one's credences in the target proposition should be responsive. Dogmatists deny that the information provided by peer disagreements plays any such role, reducing it to mere psychological information that should leave one's view on the target proposition unaffected. Non-dogmatists disagree. Non-dogmatism comes in different varieties, depending on how much weight is assigned to the counterevidence provided by peer disagreements. The best-known non-dogmatist view is the so-called Equal Weight View, according to which one ought to meet one's epistemic peer halfway, evenly splitting one's original credences and those of one's peer. Arguably, the Equal Weight View is also the most principled non-dogmatist view on the market, and throughout we will use it as our non-dogmatist paradigm.

By contrast, comparatively little attention has been devoted to the question by which process, if any, one's credences should be revised in the light of peer disagreements.[1] One might expect that this process, like other humdrum cases of adjusting one's credences to new evidence, is best conceived in terms of conditionalization: insofar as peer disagreements provide genuine counterevidence, we should update our credences by conditionalizing on it. We dub this assumption the Conditionalization Thesis.

In this paper, we will subject the Conditionalization Thesis to closer scrutiny and suggest that, once a non-dogmatist stance is taken, there are reasons to reject it. In the first part, we will motivate such a rejection on the basis of the Equal Weight View by showing that the combination of both theses cannot be upheld on a liberal interpretation of credence 1. As we shall argue, this liberal interpretation has a number of noteworthy advantages and, in addition, allows proponents of the Equal Weight View to defuse the charge of scepticism that it commonly faces. However, assuming this interpretation, proponents of the Equal Weight View – or in fact

[1] Important exceptions are Dietrich and List (2010) as well as Jehle and Fitelson (2009). See also the rich survey by Genest and Zidek (1986) concerning the pooling problem in statistics, which can be seen as a generalization of the problem of peer disagreement.

proponents of any other non-dogmatist view – have good reason to reject the Conditionalization Thesis.

In the final part of the paper, we offer an alternative construal of the process by which updating on peer disagreement should proceed. Our proposal is quite simple: peer disagreement demands a revision of prior probabilities. Instead of treating the news that a peer holds a different view as simply a further addition to our total evidence, we should conceive of it as calling into question whether we have properly weighed the force of our original evidence. Since the evidential relations we take to obtain between pieces of information is reflected in our prior probability distribution, the rational process whereby we respond to peer disagreements should involve a resetting of our original conditional probabilities. As we shall argue, this positive view is independently motivated by reflections on the nature of the counterevidence peer disagreements provide. Thus, it need not be taken to rely on the interpretation of credence 1 that we put forward. But it shows that such an interpretation can coherently be combined with the Equal Weight View.

1. Peer disagreement: the options

One's epistemic peers are subjects who are generally just as good as oneself in responding to evidence for or against a given proposition or range of propositions. Thus, they have the same epistemic powers and background knowledge, are equally smart and, in general, just as perceptive and reliable as oneself in acquiring evidence and evaluating the truth or probability of the relevant propositions in the light of that evidence (Kelly 2005: 174-75; Christensen 2011: 188-89; Kornblith 2010: 31). However, just like oneself, one's epistemic peers are not infallible, and there is no guarantee that all situations in which they are mistaken are situations in which one would likewise be mistaken, or *vice versa*. Accordingly, peerhood does not preclude the possibility that in a given situation, one disagrees with one's epistemic peers.

Let H be the relevant proposition about which one might disagree with one's epistemic peers, and let E be one's total evidence.[2] For simplicity's sake, let us assume, as it is standardly done, that E is also the total evidence of one's epistemic peers: even if someone else is, in general, as good as oneself in assessing H in the light of the kind of evidence exemplified by E, if they are better or less well informed than oneself in a given situation, there is a sense in which they do not qualify as one's peers *in that situation*.[3]

Now let D be of the form D[H,E,*z*,*x*] and express that one's epistemic peer *z* assigns credence *x* to H in the light of E. Accordingly, let P_E be one's credence function after one has learnt E but before one has learnt D, and let P_{ED} be one's credence function after one has learnt D in addition. Then all extant views in the debate about peer disagreement can be taken to suggest different values for *n* in the following schematic equation:

[2] For the purposes of this paper, evidence will be assumed to be propositional and so within the conceptual grasp of epistemic agents. This has two desired consequences: evidence can unproblematically be shared, and epistemic agents can conditionalize on the evidence they possess. We are, however, aware of the idealisation this involves.

[3] Although this assumption is standardly made (see Kelly 2005: 175; Enoch 2010: 956-57), we should note that in making it, we depart from the characterisation of epistemic peerhood given by Feldman (2011: 144). For the purposes of this paper, we will only consider those as our epistemic peers who possess the same evidence before the disagreement materialises, though in principle it would suffice to assume that none of the peers is epistemically privileged in any relevant respects *vis-à-vis* the target proposition.

(1) $P_{ED}(H) = 1/n \cdot P_E(H) + (1\text{-}1/n) \cdot x.$

For $n = 1$, we arrive at the dogmatist position according to which the credences of one's epistemic peers never rationally bear on one's own credences: learning about peer disagreement is said to have no effect on what it is rational for one to believe about the target proposition. Proponents of dogmatism include Foley (2001). For $n = 2$, by contrast, we arrive at the so-called *Equal Weight View*, according to which one's credence should meet the credence of one's epistemic peer halfway. This view has been defended by Elga (2007) and Christensen (2007a). (Note that the Equal Weight View, as here characterised, is merely a view about what the post-disagreement credences should be, and not, as some authors maintain (e.g. Fitelsen and Jehle 2009), a view that proposes an additional update rule besides conditionalization.)

On other non-dogmatist views according to which a peer's credence is given some weight but less weight than one's own, where this difference in weight is always the same, n takes a single intermediate value between 1 and 2. This kind of view has not found any supporters in the extant literature, which may well be thought unsurprising because any fixed choice of intermediate value would seem arbitrary. Note, however, that there may also be less principled views according to which the value of n may differ from case to case. This is indeed the kind of view suggested, in their respective ways, by Enoch (2010) and Kelly (2010).

Such more flexible non-dogmatist views would seem to have the advantage that they can easily accommodate cases in which epistemic peers get themselves disqualified by holding outlandish views, or cases in which they are known to be in a cognitive state that makes them poor judges of the subject matter under dispute (say because, in this particular case, they are biased for personal reasons, their general reliability notwithstanding). However, the cases of peer disagreement to which the Equal Weight View is meant to apply, and on which the debate ultimately turns, are anyway cases whose description is sufficiently hedged to exclude outlandish views or temporary cognitive impairments of such kinds (see e.g. Elga 2007: 483; Christensen 2007a: 188-89, 199). Another caveat will be that one has no personal information about the trustworthiness of the method one has applied in arriving at one's opinion that is unmatched by corresponding information about the trustworthiness of the method employed by one's peer (Christensen 2011: 8-11).[4] Although we will not go to the trouble of making all these hedge clauses explicit (and there may be more)[5], we will nonetheless assume throughout that they are in place and do not render choice between the options trivial.

In what follows, we will simply assume without argument that dogmatism gives the wrong account of how we ought to react to peer disagreements. We do believe that there are very good arguments in favour of this assumption, but we will not here undertake to provide any. For

[4] In the context of discussing this caveat, Christensen (2011: 8-9, 11-12) contends that a high degree of confidence in the reliability of one's method is correlated with a high degree of confidence in the target proposition. But this does not seem to be generally correct, as one may be highly confident that one's method reliably suggests that one's evidence mandates an intermediate credence in the target proposition.

[5] Sometimes learning that an epistemic peer agrees in taking H to be highly probable in the light of shared evidence, even if slightly less so, may make one even more confident that H holds. Thus, in cases in which the disagreement is only marginal, while there is agreement that the probability of H is very high, the significance of the latter may trump the significance of the former and turn one's peer's view into further inductive evidence for H, driving the probability of H up. It seems clear that such cases are not of primary interest in discussions about the probative force of peer *disagreements*. (Alternatively, one could classify at least some such cases as cases of peer disagreement by holding that the probability for H might indeed have to be lowered slightly. The intuition to the contrary would rather reflect the fact that our confidence in H becomes more *robust*, that is, can less easily be overridden by future evidence.)

simplicity's sake, the only non-dogmatist view we will focus on is the Equal Weight View. However, essential parts of our considerations should carry over, *mutatis mutandis*, to any non-dogmatist view.[6]

2. The Conditionalization Thesis

All non-dogmatist views maintain that one's credences should, always or occasionally, change when one finds oneself in disagreement with an epistemic peer. They all face the question by which rational process, or procedure, this change in credences is brought about.

The most straightforward answer to this question would seem to be: by conditionalization. If so, we would treat the information provided by peer disagreements just like any other ordinary kind of information and integrate it into our belief system by conditionalizing on it. After having learnt about disagreement with a peer, we should think about the relevant proposition what we thought about it antecedently to the disagreement but on the assumption that the disagreement were to occur. For instance, Elga (2007: 490) writes:

> 'Upon finding out that an advisor disagrees, your probability that you are right should equal your prior conditional probability that you would be right. Prior to what? Prior to your thinking through the disputed issue, and finding out what the advisor thinks of it. Conditional on what? On whatever you have learned about the circumstances of the disagreement.'

This we take to be a clear expression of the idea that the rational response to peer disagreements should proceed by conditionalization: our conditional credences prior to the disagreement are claimed to already anticipate its effect on our belief state should the disagreement materialise.[7] Call this the *Conditionalization Thesis*.

The Conditionalization Thesis can be seen as a natural default position given only that conditionalization is the standard updating rule in Bayesian epistemology. As far as the debate about peer disagreement is concerned, only Elga (2007) explicitly endorses the thesis in print. It

[6] According to our criteria, the so-called *Total Evidence View* proposed by Kelly (2010) qualifies as non-dogmatist and so will likewise be affected by the trilemma between non-dogmatism, liberalism about credence 1, and conditionalization to be set out below. However, this view also includes an account of how the kind of higher-order evidence exemplified by peer disagreements bears on our first-order credences – an account that conflicts with what we have to say later about the role such evidence plays; see section 6.

[7] On our interpretation of the quoted passage, Elga's clause 'prior to thinking through the disputed issue' is merely meant to ensure that having learnt E, we should not let our conditional probability of H given D, be influenced by the particular reasoning that, in the face of E, leads us to assign the unconditional probability to H that we do assign. This is motivated by the concern – sound as long as conditionalization prescribes the way in which we ought to respond to peer disagreement – that we should not demote the conflicting view we hypothesize our peer to hold simply on the basis of our own take on how probable H is made by E. There is an alternative, less straightforward, way of reading Elga's clause according to which it requires appeal to a hypothetical credence distribution that describes what we would have thought about H had we known no more than that our peer holds a conflicting opinion about H, while it brackets off what we actually think about H in the light of E. As will become evident later, this reading would move Elga's position closer to what we will ultimately suggest: adopting a hypothetical credence, instead of conditionalizing, can be seen as re-evaluating the prior probability assigned to the target proposition in the light of the original body of evidence. However, we see no clear textual evidence in Elga (2007) that would support this alternative reading.

seems that the question by which updating procedure peer disagreements might effect a change in credences has mostly remained in the background of the discussion (a clear exception is Fitelson and Jehle 2009; but Christensen 2010: 199-200, and Lasonen-Aarnio 2014: 318, also touch upon the issue.)[8]

Using the formalism introduced in the previous section, we can express this thesis schematically as follows:

(2) $P_{ED}(H) = P_E(H|D)$.

The idea that updating on peer disagreements should proceed by conditionalization is compatible with any view about how we should react to peer disagreements. What turns this idea into a specific proposal is an assumption about what our prior conditional credences should look like. Thus, (1) and (2) jointly yield

(3) $P_E(H|D) = 1/n \cdot P_E(H) + (1\text{-}1/n) \cdot x$,

where, to recall, x is the credence assigned to H in the light of E by the epistemic peer to which D relates. Consequently, we can distinguish between the possible views, just as we did before, by asking what value they assume n to take (always or in a specified range of cases).

As already advertised, we shall presume that non-dogmatism is correct and take the Equal Weight View as our non-dogmatist paradigm. It should be borne in mind in what follows that the Equal Weight View is independent from the Conditionalization Thesis. So whatever pressure might be put on the latter, this need not negatively affect the former.

3. Credence 1

The Equal Weight View is quite plausible in its realm of application, and conditionalization is the standard rule for updating our beliefs. In this section, we would like to focus on a third, but far more general issue in Bayesian epistemology: the interpretation of probability 1. We will defend a liberal interpretation and argue that it is frequently rational to assign probability 1. We proceed by first presenting some arguments in favour of this position, then supply a positive picture to underpin it, and close with briefly defending the picture against some immediate objections. The relation between liberalism about credence 1, the Equal Weight View and the Conditionalization Thesis will then concern us in the next section.

Perhaps the first thing to note is that if we update by conditionalization, the proposition we update on is subsequently assigned probability 1. By updating on E through conditionalization, we move from P to P_E, with $P_E = P(\_\,|E)$. So after the update, our new probability $P_E(E)$ equals $P(E|E) = 1$ (we renormalize our probabilities over what we have learned). Defenders of conditionalization have therefore good reason to endorse frequent assignments of probability 1 as long as they grant that we frequently learn something, i.e. provided that they are not sceptics about evidence.[9] We should also remark that some assignments of probability 1 are already

[8] As we argue in section 3 below, the criticism of standard Bayesianism to be found in Christensen (2007b) has a much more general target, i.e. the Bayesian idealisation that logical theorems should be given credence 1, and proceeds from general considerations about our fallibility.

[9] Sceptics about evidence often revert to Jeffrey conditionalization (see Jeffrey 1965: ch. 11).

forced by the standard laws of probability which require the assignment of probability 1 to all theorems of logic.[10]

Another observation is that it is hard if not impossible to avoid assignments of probability 1 to non-logical propositions in a Bayesian setting. This has to do with the fact that a probability function cannot always be *regular*.[11] Regularity means that all non-empty sets in the algebra over which a probability function is defined are assigned positive probability. Correspondingly, only the top element of the algebra – the set comprising all possibilities or the necessary proposition – is given probability 1. The problem with regularity is usually illustrated with a case in which a point-sized dart is thrown at a board with infinitely many points on it, each of which is equally likely to be hit. Regularity would require that all possible throws are assigned the same positive probability, but given the infinity of possible throws, these probabilities would then add up to more than 1. If one wants to stick to the Bayesian framework, it is hard to see how one can avoid assignments of probability 0 to propositions predicting where the dart lands, and so probability 1 to their negations (see Williamson 2007 and Hájek 2012 for further discussion).[12]

It would also be desirable if a notion of *rational acceptance* could be defined in probabilistic terms which is closed under logical consequence in the sense that rational acceptance of each of the premises entitles one to rationally accept the conclusion, upon competently deducing it from the premises. Such a notion seems to have application in the context of our ordinary reasoning. We often combine a number of premises for which we have gathered separate evidence, in order to derive something which they collectively entail and which we subsequently accept on those grounds. But as the lottery paradox illustrates, probability 1 is the only threshold $t > 0$ for which we always have it that, if two propositions have a probability of at least t, their conjunction has probability t. (For other choices of t, the probability of the conjunction can be lower than the minimal value of the conjuncts.) Given this, the only fully general way of defining a notion of rational acceptability in Bayesian terms which is closed under logical consequence is by identifying it with probability 1.[13]

For these reasons, we take it to be an attractive option to explore a liberal construal of probability 1. One such alternative takes its inspiration from Williamson's account of objective evidential probabilities and regards them as what subjective probabilities aim to match. For Williamson, evidential probabilities indicate an objective degree of justification that a body of evidence confers on a given proposition; evidential probability 1 is then interpreted as an idealized model of knowledge (Williamson 1997, 1998 and 2000: ch. 9-10). It is plausible to assume, *pace* Lewis (1996: 556), that knowledge entails belief. Accordingly, we may understand subjective probability 1 as an idealized model of the highest degree of belief implied by

[10] This may be a limitation rather than a virtue of the framework (cf. Garber 1983, Hacking 1967). But it still suggests that assignments of probability 1 to theorems of logic should be rationally permissible.

[11] It may be possible to satisfy regularity in a comparative framework for probability. See AUTHOR2 for discussion.

[12] Note, however, that regularity is unproblematic in finite settings.

[13] Apart from biting the bullet and accepting that rational acceptability is not closed under logical consequence, there have also been attempts to suggest non-general solutions to the problem by finding suitable restrictions on sets of propositions for which rational acceptability can be assumed to be closed under logical consequence. But see Douven and Williamson (2006) for critical discussion. Currently, there is an active research programme, pursued by Leitgeb (2014) and others, to revive definitions of rational acceptability according to which it does not require probability 1 but merely a suitably high probability, which will also be sufficient provided that further independent conditions are met. As it turns out, however, these proposals make rational acceptability highly context-sensitive. For instance, they imply that we may rationally accept that our lottery ticket will lose, while we cannot rationally accept this for all other lottery tickets, despite the fact that our evidence does not discriminate between them.

knowledge. Call the highest kind of belief entailed by knowledge *full belief*.[14] In this way, full belief turns out to be the subjective counterpart, or ingredient, of knowledge. And in the same sense in which it might be said that full belief aims at knowledge, subjective probabilities can be said to aim at the corresponding evidential probabilities. The resulting picture of probability 1 as full belief is very similar to the account given by Levi (1980), who takes knowledge to be the standard for *serious possibilities* and serious possibilities to be what subjective probabilities are defined over.

Just as knowledge of the conjuncts yields knowledge of the conjunction through competent deduction, rational full belief in the conjuncts entitles one to fully believe the conjunction once competently deduced. A coarse-grained model of rational full belief can therefore be expected to be closed under logical consequence. An identification of rational acceptance with rational full belief would give us a notion of rational acceptability which is not only closed under logical consequence but also frequently instantiated.

If, contrary to what a thoroughgoing subjectivism would suggest, one's credences aim to get the objective evidential probabilities right, and are to this extent subject to possible criticism,[15] why, it may be asked with hindsight, should the Equal Weight View's prescription to 'split the difference' be the least plausible? After all, if both parties to the disagreement abide by this objective standard, and so concede that either of them might on occasion fail to meet the standard, should they not now likewise concede that the objective evidential probabilities might well lie outside the interval defined by their pre-disagreement credences? If this is a possibility, splitting the difference could sometimes move one party further away from the evidential probabilities. In reply, let us first recall that regarding someone as one's epistemic peer involves taking them to have a track record that makes them as good as oneself in responding to the evidence shared. So if in a case of acknowledged disagreement with one's peer, about a given proposition A, one regards such disagreement as any reason at all to revise one's credence, as the rejection of dogmatism implies, one thereby concedes that the objective evidential probability of A and the credence one originally assigned to A possibly diverge, where such divergence is then equally likely to hold between one's own and the credence of one's peer. But this negative implication in no way undermines the positive flip-side of peerhood, namely that one's own credences and those of one's peer still remain equally good indicators of what A's objective evidential probability is, with both parties having an equally good track record that does not suddenly evaporate. Given that neither oneself nor one's peer has any independent inkling on what the objective evidential probability is (cf. Kelly 2010: 138), and given that there is no

[14] The term is supposed to contrast with *partial belief* as an elucidation of intermediate credences.

[15] Our claim that credences aim to get the evidential probabilities right might be contested on the following grounds: Even if one's credence of 0.8 in H matches the evidential probability of H, it is criticizable once it turns out later that H is false or that H was objectively very unlikely at the time of having that credence. We respond as follows: To say that credences aim to get the evidential probabilities right is not to say that they do not also aim at truth or objective likelihood. On the one hand, getting the evidential probabilities right, and so forming one's credences responsibly in the light of one's evidence, is the best way to achieve these latter aims overall and in the long run. On the other hand, one cannot coherently believe that H is true or objectively likely but that one's evidence does not support H. There is no way of aiming to get the facts right without aiming to respect the evidential probabilities. In a case in which the probability of H on one's evidence is (acknowledged to be) 0.8, it is irrational to plump for any higher or lower credence just because one aims at truth or objective likelihood. Here, the aim to respect the evidential probability and to settle for a credence of 0.8 takes precedence over the aim to get the facts right, even if it is more probable on one's evidence that one gets the facts right by fully believing H than that one gets the facts right by fully believing ¬H. In the long run, plumping for a higher credence in such cases proves detrimental to one's attempt to get the facts right.

independent indication that one rather than the other party is at fault, it accordingly seems a sound strategy to expect that both credences, though equally off the mark, get equally close to the right value.[16] There is, of course, no guarantee that this is so, i.e. the notional difference between subjective probabilities and objective evidential probabilities still leaves room for a real difference, even after such revision, where the difference may be so large that the evidential probabilities actually lay outside the interval spanned by the peers' credences. But note that the strategy of 'splitting the difference' may be rational even absent any independent positive reason for thinking that it leads to the objectively correct result in each case. The situation is comparable to cases of misleading evidence, in which updating on a piece of new information takes us further away from the truth, say by lowering our credence in a true proposition. Similarly, the credence of our peer can sometimes be misleading evidence, suggesting that we have not got the evidential probabilities right, even though our credences were in fact matching the evidential probabilities perfectly. But just like ordinary cases of misleading evidence do not suggest that there is anything wrong with the updating procedure we were following, neither do cases in which the credence of our peer is in fact far off the mark indicate that changing our priors is not the right strategy to respond to peer disagreements. Sometimes the evidence provided by a disagreeing peer will lead us astray, but it is a reasonable expectation that the credence assignments made by us and our peers tend to cluster around the true evidential probabilities (the variation in the cluster reflecting how good the peers are in evaluating the evidence), which would then make averaging a good strategy.

The present liberal conception of probability 1 contrasts with an understanding of it in terms of absolute certainty. Probability 1 is often understood as a claim to certainty beyond all possible doubts, including doubts invoking the most *recherché* sceptical scenarios. This kind of certainty frequently goes under the name of 'Cartesian certainty'. Once subjective certainty is understood as a claim to Cartesian certainty, it may well be that we should never have credence 1 in anything.[17] Considerations about the fallibility of our epistemic methods, including our logical reasoning skills, suggest that all our beliefs leave some room for doubt that they might present cases in which our epistemic methods happen to have led us astray.[18] Christensen (2007b) has

[16] Consider a case in which one has credence 0.8 in H, given E, while one's peer has credence 0.2 in H, given the very same evidence. For example, suppose that H is an inductive generalisation and that E consists of a series of observations corroborating H, while the disagreeing parties take different stands on whether these observations are representative. If it stands fifty-fifty whether they are representative, then after the disagreement manifests itself, the evidential probability of H, given the new total evidence, is certainly neither 0.2 nor 0.8. Both parties now have to deal with the doubtful representativeness of the observations made, and insofar as they both respect the evidential probability post-disagreement, it is unclear why they should not meet half way. Of course, there is no sense in which one can get the truth-value of H half right: either H is true or it is not. But this is not the present issue.

[17] The dispute between proponents of these two competing interpretations of credence 1 is not one about the meaning of 'certain' in natural language; it is rather about the question what cognitive states probability theory is best seen as modelling in quantitative terms. Insofar as proponents of the Cartesian interpretation are ready to call the state modelled by credence 1 'full belief', where full belief aims at knowledge, and agree with their opponents that epistemic possibilities should receive a positive credence, the dispute is best understood to be one about what knowledge requires. Thus, on the view subserving the Cartesian interpretation as thus conceived, what is epistemically possible is what is consistent with what one has a guarantee to know beyond any reasonable doubt, whereas on the view that subserves the interpretation we favour, what is epistemically possible is more narrowly construed as what is consistent with what one knows, including things one knows but has no guarantee to know beyond any reasonable doubt.

[18] One may know without being certain beyond all possible doubt. In cases in which one aspires to knowledge of H, acknowledging one's fallibility cannot consist in one's concession that H might be false, where H might be false just if ¬H is epistemically possible. For, to be epistemically possible is to be consistent with what one knows, and one

recently revived such a line of thought. We take these negative considerations to be sound, at least for all but the simplest of cases. However, if the laws of probability are construed as coherence constraints by which rational but imperfect epistemic subjects ought to abide, Cartesian certainty does not seem to be a good candidate interpretation of what subjective probability 1 aims at. Assigning probability 1 to a moderately complex theorem of logic would be something rational real-life subjects should rather not do. On the Cartesian interpretation, a rational subject should not be certain about a theorem even after having constructed a proof of it (Christensen 2007b: 10-12). It would thus seem that the Bayesian coherence constraints are in danger of losing any relevant connection with actual epistemic practice.

The proposed understanding of subjective probability in terms of partial belief whose upper bound is full belief is not intended to provide a complete story for all subjective aspects of our epistemic attitudes. For instance, it may be that we fully believe two propositions A and B, yet are in some sense more confident in A than in B (think of A as stating that one has hands, and of B as stating that Germany has a population of at least 80 million). So, both A and B may receive probability 1, but one has nonetheless greater confidence in A than in B – how can this be? According to the spirit of our proposal, we should start by looking at what happens at the level of knowledge. Here it seems that we can know both propositions, so that both receive evidential probability 1. Yet our justification for A is stronger than our justification for B. This may be so because, in comparison with the justification we have for B, our justification for A takes us closer to knowing that part of it not only is as strong as knowledge requires justification to be, but is actually of the right kind of pedigree to confer knowledge. Justification may be even stronger, allowing for knowledge of even higher orders. Viewed in this way, justification sufficient for knowledge of a given proposition is not necessarily the strongest possible justification obtainable for that proposition. The same could be said about full belief. Subjective probability 1 does not have to be seen as reflecting the strongest kind of possible justification.[19]

A worry about a liberal conception of probability 1 is that it could be seen to make all propositions receiving probability 1 immune to revision. The reason is the well-known fact that probability 1 cannot be undone by conditionalization: if $P(A) = 1$, then $P(A|E) = 1$ for all possible E (more on this below). But even if one does not follow Quine (1951) all the way, it is now widely held that most propositions can be revised in the light of further (possibly misleading) evidence. On our view, the problem here is ultimately not with probability 1 but with the assumption that all belief changes should be representable as applications of conditionalization (see Levi 1980: 25-26; Talbott 1991; Williamson 2000: ch. 10; cf. also

---

cannot aspire to know both H and that one does not know H. The point is familiar. Accordingly, on the liberal conception of probability 1, acknowledging one's fallibility while aspiring to knowledge of H, cannot consist in one's readiness to assign H a credence less than 1, for this would amount to the concession that ¬H is epistemically possible.

[19] If one knows on grounds independent from E that one knows H, then were one to encounter someone believing H to a lesser degree in the light of E, one might be permitted to discount their view as inadequate. But in such situations, the other party will lack evidence one possesses and so will not count as a peer. If one's claim to second-order knowledge essentially relies on the assumption, correct as it might be, that one has gauged the force of E right, where E is all the evidence one possesses, then the finding that one's peer gauges it differently should give one pause for thought. Thus, claims to second-order knowledge are no less subject to reassessment in the light of peer disagreement. In general, if one gets the evidential probability right, one need not be privy to the fact that one does: it may be opaque what the evidential force of one's evidence is, even if one measures it right. It is an interesting question how stronger kinds of justification, approximating knowledge of higher orders, could be modelled at the level of belief. We conjecture that the framework of ranking functions is a good starting point to collect the necessary resources (see Spohn 2012 for a recent development).

Titelbaum 2013). Perhaps the most straightforward cases are memory losses: sometimes our body of evidence does not expand but rather contracts because we forget something. As we shall argue in more detail in the following sections, peer disagreement can provide another instance of failures of conditionalization. For now it should suffice to say that we do not take probability 1 to be a mark of immunity to revision or intransigence.

Finally, let us also mention one controversial aspect of probability 1, only to set it aside in what follows. On standard decision-theoretical assumptions, we should evaluate fair bets in terms of their expected value and prefer a bet over others if it has greater expected value. This seems to have counterintuitive consequences when it comes to beliefs held with probability 1, because any odds appear favourable once we are certain about the relevant outcome. As a consequence, we would have to bet our life against a penny, which seems absurd. Recently, there has been a renewed interest in this problem in connection with the claim, made by Williamson (1998 and 2000: ch. 9-10), that knowledge has evidential probability 1 (see Greco 2013; Kaplan 2009). There are various possible reactions to the problem (Greco 2013 gives a critical overview). For instance, one could break the allegedly tight connection between decision and probability by arguing that higher stakes require more epistemic care than lower stakes, and thereby demand knowledge of higher order (see Williamson 2005 and 2009b).[20]

4. A trilemma

In the preceding section, we argued that there is good reason to adopt an interpretation of certainty that leaves room for the frequent rationality of assigning probability 1. However, it can easily be seen that the Equal Weight View, the Conditionalization Thesis and the liberalist claim just defended – that certain propositions, including non-logical ones, might rationally be assigned probability 1 – will form an inconsistent trias in those cases in which such propositions include propositions which are the possible objects of peer disagreement.

Let H be such a proposition and let D again encapsulate the information that an epistemic peer disagrees with one's probability assignment to H, where D itself is assigned a non-zero probability. Note that D may be assigned a non-zero probability consistently with H's receiving probability 1, insofar as one may well think in advance that an epistemic peer who disagreed would make a mistake, her status as a peer notwithstanding. Given the Equal Weight View, D not only provides biographical information about one's peer, but constitutes genuine counterevidence. Accordingly, on this view, one should still lower one's credence in H upon learning D. Given the Conditionalization Thesis, this revision of one's credence should already be reflected in one's previous conditional probability of H given D. But now, insofar as H was originally assigned probability 1, this conditional probability had likewise to be 1. For, by the definition of conditional probabilities, if $P(D) > 0$, then $P(H|D) = P(H \,\&\, D)/P(D)$, and by the laws of probability, if $P(H) = 1$, then $P(H \,\&\, D) = P(D)$. By conditionalization, $P(H|D) = P_D(H) = 1$. Consequently, one would not be able to comply with the demands of the Equal Weight View in such a case, because *ex hypothesi*, one's peer assigned H a probability *m* less than 1 and so, on this view, $P_D(H)$ should equal $(m + 1)/2$.[21]

[20] For more on the compatibility of credence 1 with refusals to bet one's life against a penny, see AUTHOR1.

[21] It might be thought that the Equal Weight View does not sit well with the kind of liberalism about credence 1 that we favour – and this quite independently from the Conditionalization Thesis. For, once it is assumed, with Williamson (2000), that whatever we know is part of our evidence, it would seem that evidence E can license

The current problem for conditionalization is not just the problem that we cannot update on a proposition which we have antecedently assigned zero probability, for D is assigned a positive probability, making the conditional probability well-defined. Rather, the problem is an instance of the problem that probability 1 cannot be undone by conditionalization even if a sufficiently strong defeater becomes known. Despite being a fairly well-known problem, it currently does not have a particularly rich solution space (see, however, Titelbaum 2013). We think that the case of peer disagreement is instructive because it helps us to identify a specific way of solving the problem for a certain class of particularly interesting cases.

One might be tempted to resolve the aforementioned conflict by denying that those propositions which can rationally be assigned probability 1 can ever be the objects of peer disagreement. It might then seem that all three claims can after all coherently be endorsed. But on reflection, this attempt at harmonization fails. For, the possibility of peer disagreement is ubiquitous, in the sense that almost no non-logical proposition can ever be said to be immune to possibly conflicting probability assignments by epistemic peers. As debates in the philosophy of logic show, even logical propositions are candidate objects of peer disagreement. Thus, ultimately, the combination of the Equal Weight View with the Conditionalization Thesis could be upheld only at the expense of bereaving the third claim, that some propositions might rationally be assigned probability 1, of much of its interest.[22]

From the very beginning, we set aside dogmatist accounts that deny peer disagreements the status of genuine counterevidence to which one's probability assignments should be responsive. So, for the purposes of the present discussion, non-dogmatism about peer disagreement is assumed as non-negotiable. Since non-dogmatist rivals to the Equal Weight View will be subject

---

assignment of credence 1 to H (in the sense that the evidential probability of H given E is 1) only if E entails H – in which case any disagreement about the bearing E has on H would betray probabilistic incoherence on the part of one of the disagreeing parties. Accordingly, as long as we also follow standard Bayesianism in supposing that peers abide by the ideal of probabilistic coherence, there would be no cases in which one peer correctly assigns credence 1, while the other does not. In response, let us first note that much of the extant literature on peer disagreement would seem to presuppose a more restrictive notion of evidence than the one Williamson assumes – a notion according to which what we know is *not* automatically part of our evidence. Secondly, even if Williamson's equation of knowledge with evidence was accepted, the connection between conferral of evidential probability 1 and entailment still crucially depends on the assumption of regularity – an assumption that Williamson (2007) and Hájek (2012) provide reasons to reject.

[22] But why shouldn't one be in a position to demote anyone from the status of being a peer who has the very same evidence and yet fails to fully believe what one takes as known in the light of that evidence? This is the kind of suggestion dogmatists are notoriously prone to making; and it also has an evident bearing on Kripke's puzzle that we comment upon in footnote 28. Let us here briefly say something in response, if only to dispel the false impression that appeal to credence 1 does any extra work. First, since one may know a proposition without knowing that one knows it, taking oneself to know a proposition does not imply taking oneself to have a guarantee that one knows it, and hence neither taking oneself to have a guarantee that in assigning it credence 1, one has properly gauged the force of one's evidence. Secondly, to the extent that, in general, when assigning probabilities to propositions in the light of one's evidence, one aims to get their evidential probabilities right, and insofar as the track records that make one an equal peer are records of success in pursuit of this aim, there is no reason to expect that in cases in which one assigns credence 1 to a given proposition one pretends to be on safer ground as far as one's assessment of the evidence is concerned than in cases in which one assigns an intermediate credence. If in the latter type of case, there is reason not to dismiss the divergent credence of someone with an equally good track record as a sure sign of their being mistaken, then so there is in the former type of case (cf. AUTHOR1). As far as we can see, disagreements about which intermediate credences to assign are continuous with disagreements in which one party assigns credence 1.

to the very same problem diagnosed here, we might just as well continue treating the Equal Weight View as our non-dogmatist paradigm.[23]

What other options are there for resolving the diagnosed conflict? That probabilities 1 cannot be undone by updating on one's evidence is a familiar, if controversial, feature of conditionalization. In this respect, updating on the evidence provided by peer disagreements is just a special case. This general observation might already be taken to be sufficient ground for thinking that non-logical propositions should never be assigned probability 1, as we should keep an open mind about them and allow for our views to be corrected by future evidence. Such is indeed a prominent reaction. For instance, this thought can be taken to yield one motivation for Jeffrey's rule that allows for updates on uncertain evidence, so that propositions encapsulating one's evidence do not automatically receive credence 1 (see Jeffrey 1965: ch. 11).

It should be noted, though, that Jeffrey's rule is a generalized version of conditionalization that still allows for assignments of probability 1 as a limiting case. In this limiting case, Jeffrey's rule delivers the same verdicts as conditionalization and so gives rise to the very same kind of conflict diagnosed above. For, given that there is no principled distinction between non-logical propositions that are, and those that are not, apt to encode one's evidence, the claim that non-logical propositions may rationally be assigned probability 1 naturally extends to propositions encoding one's evidence. In other words, for all that Jeffrey's rule dictates, even if not all evidence is certain, not all evidence is uncertain. Advertising Jeffrey's rule as an improvement upon standard conditionalization is certainly not enough to air the view that one ought never to assign credence 1 to non-logical propositions.

It remains unclear, however, how denying that non-logical propositions could ever rationally be assigned probability 1 can be motivated on independent grounds – that is, without simply taking it for granted that conditionalization is sacrosanct. Such independent grounds would seem to become readily available once one adopts the first of the two interpretations of subjective certainty discussed in the previous section. Thus, it should be conceded by all parties that aspiration to Cartesian certainty about a given proposition is incompatible with one's concession that future evidence, and be it of the type D exemplifies, should ever mandate revising one's belief in that proposition. After all, on a natural understanding, a claim to Cartesian certainty involves a claim to incorrigibility. Thus, interpreting probability 1 in these terms indeed supplies an independent reason for denying that it could ever rationally be assigned to propositions to which the Equal Weight View applies. The ubiquitousness of peer disagreement will no longer tell against the implied restriction of the latter view's range of application, as rational aspirations to Cartesian certainty are correspondingly rare.

But things look importantly different once we consider the alternative interpretation of subjective certainty in terms of full belief. If we understand probability 1 in this less demanding way, there is no evident reason to think that non-logical propositions cannot frequently be assigned probability 1, and rationally so. And yet, for all that has been said, one's full belief in a given proposition would seem to be perfectly compatible with one's concession that, for all one knows, future evidence, once available, might rationally mandate revising one's belief (even if it should be misleading).[24]

---

[23] Even on Kelly's Total Evidence View, at least in some cases, upon learning information like D one should 'temper one's confidence' in the target proposition H (Kelly 2010: 138-39). There is no reason why such cases cannot include cases in which one's original degree of credence in H is 1.

[24] The latter suggestion is in conflict with van Fraassen's *special reflection principle* (see van Fraassen 1984 and 1995; see also Weisberg 2005 for a comparison with conditionalization). This principle is controversial for reasons

Conditionalization treats acquisition of evidence as cumulative. It is for this reason that it cannot cater for cases in which parts of one's original evidence are lost and the impoverished body of evidence rationally mandates a different probability assignment to the target proposition. Paradigmatic such cases are cases in which evidence is simply forgotten (Williamson 2000: 219; Talbott 1991; Titelbaum 2013). Thus, if part of one's original evidence, on whose basis P(H|A) is assigned probability 1, is forgotten in the process of acquiring A, $P_A(H)$ may turn out to be distinct from 1. Whatever its details, the most natural way to accommodate memory loss will also allow for probability 1 to be undone (cf. Titelbaum 2013).[25] Accordingly, appeal to conditionalization cannot provide a legitimate basis for denying the general claim that assignments of probability 1 may be rational even when one's future evidential situation might, for all one knows, mandate a lower credence.

Promoting the failure of conditionalization comes at a well-known cost. For one must then face up to the challenge that epistemic agents become vulnerable to so-called diachronic Dutch-books. Whether this challenge can be defused by denying the close link between probabilities and betting-behaviour which Dutch-books exploit, or whether one should rather bite the bullet and concede that epistemic rationality sometimes makes such vulnerability unavoidable, is a matter of some controversy (cf. Christensen 1991; Williamson 2000: 219; AUTHOR1) – a controversy which we cannot take up here.

The reasons we have given do not in themselves constitute an argument for the claim that probability 1 can rationally be undone in the light of certain counterevidence, such as, for example, peer disagreement. But they remove one stumbling block on the way towards devising such an argument: contrary to what one might have thought, Bayesian conditionalization cannot claim to capture all there is to say about rational belief revision anyway. The diagnosed conflict with the other two claims of our triad, together with the reasons in favour of those two claims, may then be taken to provide some basis for the thought that conditionalization might after all not be the right principle of belief revision to look to in the context of peer disagreement.

The remaining option accordingly is to reject the idea that update by conditionalization is the appropriate procedure to regulate our rational response to peer disagreements.[26] As we shall argue in section 6, such rejection can independently be motivated, given the special nature of the counterevidence peer disagreements provide. In the next section, we wish to highlight a theoretical virtue that rejection of this idea should be said to have for proponents of the Equal Weight View.

---

having to do *inter alia* with information loss or a deterioration in our standards of assessment (see Talbott 1991 and Briggs 2009 for discussion). Now, an envisaged change in our prior probabilities due to peer disagreement is a change in our standards of assessment, though it need not always be taken to be a change for the worse, as we can be agnostic about whether or not one of our peers might work with more adequate prior probabilities. Still, the predicted type of counterexample to the reflection principle is in the immediate vicinity of the counterexamples we are already familiar with.

[25] Recall that not even Jeffrey's rule can undo probability 1.

[26] Fitelson and Jehle (2009) show that certain general principles of updating, strong enough to capture the Equal Weight View, yield violations of conditionalization. However, they take this as a reason to weaken the relevant principles, so that the Equal Weight View is only being approximated. Shogenji (2007) also contains a counterexample to conditionalization (together with an extensive overview of possible reactions). However, Shogenji works with a much stronger conception of what the Equal Weight View amounts to – too strong a conception, we think. We will further discuss Fitelsen and Jehle (2009) in section 8 below.

5. Anti-scepticism

The Equal Weight View does forbid credence 1 in situations in which evidence such as D has become available: disagreement implies that at least one of the disagreeing parties assigns H a probability *m* such that $(m+1)/2 < 1$. By contrast, the Equal Weight View as such makes no demands on the probability assigned to H before peer disagreements manifest themselves. If conditionalization fails, or fails to apply, that probability may accordingly well be 1.

Proponents of the Equal Weight View should welcome this result. Typically, the Equal Weight View is accused of having unacceptably far-reaching sceptical consequences. One way to understand this accusation makes it depend on the assumption that actual manifestations of peer disagreement abound. However, it is unclear whether this assumption is justified. Recall that in order to properly perceive a situation as a situation in which one disagrees with a peer about the probative force of one's evidence, one not only must be entitled to consider the disagreeing party as a peer, but also be entitled to think that they are in possession of exactly the same evidence and are as good as oneself in evaluating that particular body of evidence, where the latter is not already guaranteed by their good overall track record. Such cases may well be less frequent than current discussion tends to make one believe. More often than not it is an open question whether one's peer does not after all possess more or less information than oneself, and cases may be very few where this question can confidently be answered in the negative.

This does not entail, however, that it should not frequently be considered a genuine possibility that an epistemic peer might disagree about the support provided by the very evidence in one's possession. So proponents of the Equal Weight View should at least concede that the *possibility* of peer disagreement abounds. As long as the Conditionalization Thesis remains in place, this concession will force them to conclude that assignments of probability 1 are rarely, if ever, rational. So an alternative way to understand the charge of scepticism makes it depend on the inference that takes us from this concession to the conclusion that the Equal Weight View mandates less than full belief in areas where we would typically take ourselves to be entitled to assignments of probability 1.

As we have seen, it is true by virtue of the definition of conditional probabilities and the laws of probability, that if $P(D) > 0$ and $P(H) = 1$, then $P(H|D) = 1$. So it does follow from the fact that an epistemic agent assigns probability 1 to a given proposition H, that she must, at that stage, presume that D would provide misleading evidence, were it to become manifest. But this alone does not show such an assignment to be in conflict with the Equal Weight View: on this view, even one's epistemic peers occasionally get matters wrong, and one might on occasion rationally take oneself to be entitled to think they would get matters wrong, were they to disagree. Only if one thought that $P(H|D)$'s being 1 implied that, faced with D, the epistemic agent would be bound to retain his credence in H, might one take this to suggest any incompatibility between the Equal Weight View and our widespread entitlement to full belief (cf. Christensen 2010: 199-200; AUTHOR1). Yet, this diachronic connection will be mandatory only as long as the Conditionalization Thesis remains in force. Accordingly, if probability assignments are no longer bound by the demands imposed by conditionalization, this second charge of scepticism can easily be averted – which, as we submit, would make acceptance of the Equal Weight View more palatable than it might otherwise seem.

As already adumbrated, in the following section we shall argue that there is independent reason to think that our rational response to peer disagreement is best seen not to proceed by

conditionalization. Our argument will turn on the special kind of counterevidence peer disagreements provide.[27]

6. Peer disagreement demands a revision of priors

So far we have given some initial reasons to think that the culprit in our trilemma is the Conditionalization Thesis rather than the Equal Weight View or the claim that it is often rational to assign credence 1. We do not want to pretend that those reasons are conclusive, but we hope they are strong enough to make exploring a possible alternative worthwhile. In this section, we present our preferred view: our response to peer disagreement should primarily be an adjustment of our prior probabilities, as peer disagreement concerns the way we have weighed the evidence.[28] Thus, a rational response to peer disagreement need not proceed via conditionalization but could equally well involve a revision of our prior probabilities. This view has the evident advantage that it permits us to combine the Equal Weight View with a liberal construal of credence 1, which in turn allows rational credences with that value to be widespread. We will now argue that reflection on what peer disagreement is about provides independent reasons for this view – that is, reasons that can be appreciated independently from the trilemma we set out in the preceding sections.

Normally, a piece of counterevidence is silent on what it is rational to believe, all things considered, before it itself becomes available. Thus, for example, if one's original evidence

[27] We will not take issue with the idealisation, inherent in standard Bayesianism, that logical theorems should receive credence 1. For one thing, we remain unconvinced by the fallibility-based reasoning that leads Christensen (2007b) to conclude that probabilistic coherence cannot even serve as a rational ideal (see AUTHOR1 for discussion). For another, our criticism of standard Bayesianism will survive even if this ideal remains in play, which should also dispel the impression that the problem we tackle is simply an instance of a much more general problem which Christensen (2007b) identified, yet which we never fully address. By the same token, we throughout work with a version of the Equal Weight View according to which peer disagreements over logical theorems are outside its scope. This might prompt the following complaint: if such disagreements are exempted from its scope, this must be because, on that view, there is a unique rational credence to be assigned to logical theorems. But then it would seem that, obversely, disagreements within its scope are those concerning propositions for which such uniqueness claim fails to hold. And yet, the view's commendation to split the difference proves entirely unmotivated unless there is a unique rational credence the target proposition merits. For why else should one move one's credence towards that of one's peer? The complaint would seem to equivocate between two senses of 'rational': it is one thing to be rational in meeting the standards of probabilistic coherence, it is another to be rational in meeting the standards set by the objective evidential probabilities. The uniqueness that the Equal Weight View requires relates to the latter, so that it would seem that the charge of invidious discrimination proves unfounded. However, one might nonetheless worry that the distinction between the kind of idealisation inherent in standard Bayesianism and the idealisation that would assume the agent to get all the evidential probabilities of propositions right, be they logical or non-logical in character, is of little epistemological interest. We disagree. In order to properly assess what follows logically from what, fewer capacities are required than those required in order to properly assess what supports what. A perfect theorem prover will in general be a bad tool in order to figure out to what degree the empirical evidence supports the conclusion of an ampliative inference. For the latter, the off-line use of perceptual capacities might be required, where there is no guarantee that the tacit assumptions made in its course correspond to propositions to which the agent has assigned any probabilities. Having said that, let us emphasize that we ackowledge that peer disagreements about logical matters are possible, and so consider the fact that they cannot be modelled in standard Bayesian frameworks as a limitation rather than a virtue of these frameworks.

[28] After submission of this article, we learnt that Brössel and Eder (2014) defend a similar type of answer to the question of what kind of revision peer disagreements mandate. Our own ideas have been developed in complete independence from theirs.

consists of so many observed white swans and no observed black one, the recalcitrant datum of a black swan-like creature, though overruling the force of one's original evidence to justify the hypothesis that all swans are white, has no bearing at all on whether it was rational to accept that hypothesis on the basis of one's original evidence, as long as the latter was all there was to consider. Similarly, if one's original evidence consists of certain observations of a given scene, the finding that atmospheric conditions distorted one's view, though undercutting the probative force of one's observations, has no bearing at all on whether it was rational to trust these observations as long as they were all the information one had to go on.

By contrast, disagreements with peers, apprised of the same original evidence, are pieces of counterevidence that call into question whether one's original evidence really justified one's hypothesis to the degree to which one took it to do so, even in situations in which that evidence was one's total evidence. As Lasonen-Aarnio puts it, such higher-order evidence

> 'has a retrospective aspect, providing a subject with evidence that her belief was never rational, reasonable, or justified to start out with. It's not just that now, once I get [such higher-order evidence], I am no longer justified in believing what I did. I acquire evidence that I was never justified to start out with' (Lasonen-Aarnio 2014: 317).

Thus counterevidence supplied by peer disagreements is of a peculiar kind. Peer disagreements cast rational doubt on whether one's response to that original evidence was properly calibrated. The difference in credence between two peers typically reflects a difference in how likely they think the target proposition is made by the evidence. Given this diagnosis, it would be most natural to model the belief change induced by peer disagreement as resulting primarily from a revision of our assessment of the force of the evidence.[29] This contrasts with the idea of describing it in terms of conditionalization. For, contrary to what the Conditionalization Thesis suggests, the information provided by peer disagreements would not then be treated as a simple addition to one's evidence, to the effect that one's new total evidence mandated a credence in the target proposition different from the one mandated by one's previous total evidence. Rather, we should then conceive of the stock of our evidence bearing on the target proposition as remaining basically the same, and so as being itself unaffected by the disagreement, and rather change the way in which we evaluate its rational bearing on the target proposition.

For instance, suppose that two social scientists who treat each other as epistemic peers come to different conclusions about whether a given body of data suggests that street lights have a positive effect on the crime rate in a given region. On the present view, their reaction to the disagreement should primarily consist in revising their take on the relevant body of data. If we add the Equal Weight View to this picture, the revision of their priors should lead them to evenly split their conditional credences in the target proposition given the body of data. Intuitively, this seems to reflect what their disagreement is about: it suggests that the subject which antecedently had a higher credence might have overrated the force of the evidence, while her peer might have underrated it.

[29] Of course, in addition to being evidence against our priors, peer disagreement also provides purely factual information about the mental state of our peer, namely that she has currently such-and-such a credence in the proposition under consideration. Thus the evidential import of peer disagreement is in fact twofold. Kelly's label 'psychological evidence' for the evidence provided by peer disagreement is therefore slightly misleading (Kelly 2010: 128). We will come back to this in section 7 below.

This diagnosis has some bearing on the current debate about higher-order evidence. For instance, Kelly (2010) objects to the Equal Weight View on the grounds that it lets the higher-order evidence for conflicts of opinion with one's peer 'completely swamp' one's original first-order evidence that is now given 'no weight' at all and 'counts for' nothing, although it is not lost in the learning process either (Kelly 2010: 124, 130, 133, 141; cf. also Kelly 2005: 190; Weatherson unpublished). According to Kelly, '[i]t is a weakness of the Equal Weight View that it assimilates cases in which one does have access to the original evidence to cases in which one does not' (Kelly 2010: 124). Even those who harbour sympathies for the Equal Weight View sometimes write in ways that would suggest this. Thus, for example, Christensen (2010) writes that when learning higher-order evidence, and trying to do it justice, 'I must in some sense, and to at least some extent, *put aside* or *bracket* my original reasons […]. In a sense, I am barred from giving a certain part of my evidence its due' (Christensen 2010: 195; emphases in the original). These considerations strike us as odd on several counts.

The conflicting opinion of one's peer challenges the appropriateness of one's own response to the original evidence. Assuming non-dogmatism, one should accordingly be responsive to this challenge and feel called upon to assuage the doubt it instils.[30] Yet, as Kelly (2010: 159) concedes, even if one got the probative force of the original evidence right, the cognitive process by which one came to appreciate the force of one's evidence, and formed one's opinion accordingly, typically neither depends nor makes available higher-order evidence that one has gauged that force correctly. Now even if that very process should produce higher-order knowledge about the force of one's evidence – as Kelly (2010: 156) contends and contrary to what qualms about the principled failure of the KK-principle suggest – such knowledge would be put just as much in doubt as one's first-order credence insofar as it would then be the outcome of the very same process.[31] It would accordingly be of no avail in order to address the challenge. For, as Kelly (2010: 138) is also ready to admit, 'one does not occupy the God's-eye point of view with respect to the question of who has evaluated the evidence correctly and who has not'; and so even higher-order evidence that is in fact misleading is evidence one is rationally bound to come to terms with (Kelly 2005: 186; Kelly 2010: 137; Christensen 2010: 191). Accordingly, the challenge posed by the peer's conflicting opinion still stands.

In such cases, the Equal Weight View counsels to split the difference. But this ruling does not imply that one's original evidence counts for nothing. To begin with, the disagreement concerns the question of what credence H merits given evidence E, and as such it would have no relevance at all for one's present opinion on H, if one did not actually possessed E: so E is required, in order for the peer disagreement to unfold its claimed potential. Furthermore note that as long as the original credences involved do not add up to 1, the advice to split the difference already presupposes that the original evidence counts for something: for otherwise abstention would be the only reasonable response. Lastly, the Equal Weight View does not bar one from giving one's evidence its due either. For, upon first learning about the conflict of opinion, one enters into a state of unclarity as to *what exactly that due is*. The view then counsels that, in the light of such

[30] Note that the doubt is instilled in one's own mind, and neither the need nor the means to assuage it should be conflated with the need or means to persuade one's peer (cf. Kelly 2010: 171-72, on the 'dialectical conception of evidence').

[31] Kelly suggests that acquisition of first-order knowledge *generates* knowledge about the force of one's evidence. On any mildly externalist view, one may come to know a given proposition on the basis of one's evidence, without being in a position to know that one's evidence has the right kind of pedigree. Though typically presented as a reason against the KK-principle, it also counts against Kelly's claim, as long as the force of one's evidence changes with its pedigree.

unclarity, splitting the difference is the reasonable way to respond, given that one's own credentials match those of one's peer. This would not make much sense if the probative force of the original evidence *as one originally conceived it* was put aside. Hence far from swamping one's original evidence, the counterevidence provided by peer disagreements works as a catalyst for the reassessment of the relation between the target proposition and one's original evidence which, as far as the evaluation of that proposition goes, still counts for everything.[32]

The present suggestion translates into the Bayesian framework as follows. Our current credence in a proposition A, $P_E(A)$, is a product of the evidence E we possess and our initial conditional probabilities. So, our present credences are given by our initial probabilities, conditional on the evidence we possess: $P_E(A) = P(A|E)$. In this framework, the way we assess the evidence is encoded in the initial (or prior) probability distribution P. The initial probabilities reflect our take on the evidential relations between the body of evidence E and proposition A. Our final credences then result from applying our assessment of the evidence, as captured by the initial conditional probabilities, to the evidence at hand. What we suggest is that peer disagreements should induce a revision of our initial conditional probabilities. Rather than taking peer disagreement to provide a further piece of evidence, assessed by the same prior probability function, the belief change mandated by disclosed peer disagreement should typically come about by revising our initial conditional probabilities, leading to a reassessment of the probative force of the original body of evidence. To put it in the form of a slogan: peer disagreement demands a revision of prior probabilities.[33]

In line with what is here being suggested, when we speak of 'revising prior probabilities' we do not merely intend the post-disagreement revision of pre-disagreement credences in the target proposition, and potentially in propositions entailing it – something that the combination of the Equal Weight View and the Conditionalization Thesis likewise implies. Rather, we here intend the post-disagreement revision of the pre-disagreement credences conditional on the shared evidence, and hence of the features of the initial probability distribution – with which we started at the beginning of our epistemic lives, as it were – that gave rise to those conditional credences.

---

[32] To further elaborate this point: After the disagreement manifests itself, one indeed retains the evidence one had before. But it is not as if, after the disagreement manifests itself, the evidence can still rationally be taken to provide reasons for setting the credence where one set it before. Suppose one has made *n* observations from *m* samples corroborating a hypothesis H, and on that basis comes to assign H probability 0.8, thereby taking the *m* samples to be representative and the *n* observations to be statistically significant. Then someone comes along and casts doubt on whether the samples are representative, therefore giving H credence 0.2. One cannot then reason as follows: 'I have made my *n* observations from my *m* samples – nothing takes them away from my evidential basis – so they still speak in favour of setting my credence in H to 0.8. It's just that I now also have to accommodate the disagreement data. So my credence in H should be slightly less, because I have one recalcitrant datum more. By comparison to my *n* obervations, where *n* is a rather large number, one recalcitrant datum does not count for much, at least not as much as to meet my peer half way.' This line of reasoning wrongly assumes that one can retain one's confidence in how probable one's observations make H, when taken on their own. We argue that this is mistaken: the disagreement with one's peer casts doubt on whether one's pre-disagreement evidence really has the evidential import when taken on its own. If this is right, then one cannot simply offset one set of data against the other. Now one may even agree with one's peer that if the samples aren't representative, the credence should lie where one's peer sets it, i.e. at 0.2. But we are peers. So my take on what's representative has a 50% chance of being right and their take on what's representative also has a 50% chance of being right. It is therefore sensible neither to stick to one's guns, nor to offset the disagreement data against the *n* observations, nor to defer, but rather to take the middle ground.

[33] If we are right in contending that peer disagreement calls for a revision of one's prior conditional probabilities, then this not only tells against standard updating by conditionalization, but also against any other proposal to respond to such disagreements, say by an inference to the best explanation. However, for want of space, and given the focus of the present paper, we cannot discuss such alternative proposals.

(See section 7 for more details.) For, if our diagnosis is correct, peer disagreement challenges one's earlier take on the probative force of that evidence when that evidence was all the evidence there was and so even before the disagreement became manifest.

Let us stress that this suggestion is compatible with standard views on what our final credences in the target proposition should look like. In particular, it is compatible with the Equal Weight View – at least as long as this view is construed exclusively as a view about the credences we finally ought to arrive at, and not as a view about what our conditional credences should look like before we learn about the disagreement (see section 2). What is new on the present picture is the description of the mechanism by which we change our beliefs in the light of peer disagreements. Assuming the Equal Weight View, we would split our credences in the target proposition A because we split the relevant part of our old initial conditional probabilities.

In particular, the idea would be that we adjust our initial probabilities in such a way that our new initial conditional probability of A given E is evenly split between our original one and that of our peer. Formally: $P_{new}(A|E) = ½ \cdot P_{old}(A|E) + ½ \cdot P_{peer}(A|E)$. At the end of the process, our revised credence in A is the result of applying the adjusted initial conditional probabilities to our original body of evidence E.

If one wishes to account for the belief change induced by peer disagreements in terms of conditionalization, one needs to build the effect of such disagreements into the conditional probabilities the subject has before the disagreements occur. While the present view can allow for $P_{old}(A|E \,\&\, D) = P_{old}(A|E)$, where P_{old} represents our initial credences before the relevant disagreement D becomes manifest, this cannot be tolerated if the subject is to learn from that disagreement by conditionalization and is expected to change her credences – say, according to the Equal Weight View. As a consequence, the Conditionalization Thesis considers certain beliefs, held before the disagreement becomes manifest, as being irrational which, on the present view, are perfectly in order. As we have argued, it is sometimes fine (a) to stick to one's credence in the target proposition even under the assumption that a peer disagrees (thinking that in this case she would make a mistake), but (b) change one's credence once the disagreement actually materialises. As previously indicated, such a situation is bound to occur when we are antecedently certain about the target proposition A, for then we have $P_{old}(A|E \,\&\, D) = P_{old}(A|E) = 1$ already by the laws of probability.[34] On the present view, changing our priors, so that E no longer makes A certain, is a way around this issue.[35]

[34] In section 3, we argued for the liberal interpretation of credence 1 according to which it is the subjective correlate of knowledge. Provided that this interpretation is correct, the cases envisaged here bear a striking resemblance to the cases invoked by Harman (1973) in his discussion of Kripke's puzzle of dogmatism, in which the right to claim knowledge is lost (cf. also Ginet 1980). According to Harman, we may think, based on our knowledge of *p,* that potential counterevidence to *p* will be misleading, but once that counterevidence presents itself, we may no longer be in a position to dismiss it as misleading, for we may have lost our knowledge of *p* in the process. To admit, Kripke's puzzle has also received conditionalization-friendly treatment, but only at the cost of rejecting the liberal interpretation (cf. Sorensen 1988). Quite independently, however, there is an intuitive contrast between what one may legitimately think about potential counterevidence in the light of one's actual evidence, on the one hand – e.g. that such counterevidence would be misleading and hence ought to be ignored – and what one may legitimately think about that counterevidence once it is actually added to one's evidence, on the other (cf. Christensen 2010: 199-200). In an appendix to his (2011), Kripke argues that Harman's solution fails to resolve the puzzle, on the grounds that the puzzle is essentially one 'about a resolution *made in advance* to ignore certain types of evidence' and so 'not to get into […] a situation' in which one may lose one's knowledge 'and hence no longer know that the counterevidence is misleading' (Kripke 2011: 48-49; emphasis in the original). Kripke's puzzle is a puzzle and would certainly demand a more thorough discussion than the scope of the present paper permits. However, let us make the following observations that apply to the case of present interest: Even if one has an antecedent reason to

7. Modelling revisions of initial probabilities

So far, we have not yet said anything about how a revision of initial probabilities might work. The discussion in the previous section has made it clear that the revision should be subject to the constraint that the revised initial probabilities assign to the conditional probability of our target hypothesis given our evidence a certain value x, where x results from moving our conditional probability towards that of our peer. Thus, the revision of priors we are looking for is a form of *constrained belief revision.* Revision rules of this kind have been studied in the literature, notably by Skyrms (1987), whose model we shall use to describe how initial probabilities might be revised (our exposition draws on Joyce 1999, 184ff.).[36] The basic idea is that given a probability function P which is to be revised in order to satisfy a constraint C, we go to a probability function P′ which satisfies C and is in some sense the most similar function with this property. In other words, revising a probability function according to a constraint C is to minimally revise P so that it satisfies C.

We shall model a belief state as a pair (P, E) of an initial probability function P and a body of evidence E. Our credences in an arbitrary proposition A according to such a belief state will then be our conditional probability P(A|E). This reflects the Bayesian assumption that our present credences come from an initial credence distribution by conditionalizing on our evidence. In addition, by representing a belief state as a pair (P, E), we assume – thereby slightly going beyond standard Bayesian assumptions – that an initial probability function can still be recovered

---

systematically avoid getting oneself into a position in which one is faced with the kind of sustained peer disagreement that forces revision, it is hard to see how to make this a reason for one's actions. The evidence provided by peer disagreement is evidence that turns out to be counterevidence only when it is too late to ignore it, by which time it is undetectably misleading, if it is misleading. Upon meeting one's peer, one is reasonably confident that she will either agree with one or that one can win her over after some argument. Thus one has no antecedent reason to systematically avoid encounters with one's peer, because one has no antecedent reason to expect that she will not come over to one's side, even if it takes a little argument to achieve this. When it turns out that she will not, it is too late to ignore her testimony, but at that time, one is no longer able to recognise that her testimony is misleading, even if it is, because by that time one's confidence in the original hypothesis should be shaken.

[35] We should note that changing our priors in this way is not possible when E logically entails A, because any probability function P then has P(A|E) = 1, by the laws of probability. As mentioned earlier, certain situations may lead one to delete parts of what one originally counted as evidence in one's possession – a kind of 'contraction' of evidence that conditionalization is notoriously unable to accommodate (cf. Gärdenfors 1988: ch. 5; Williamson 2000: ch. 10). Thus cases in which one has credence 1 in a given proposition and then finds oneself in disagreement with someone else on the matter, may often be cases in which some of the evidence one took oneself to have is undermined and so should be discarded. This is provably so when the evidence one took oneself to have entails the target proposition. However, it is difficult to see how discarding part of one's evidence in response to disagreement with another party, and adjusting one's credence in the target proposition accordingly, can be rational if the disagreeing party arrived at her distinct credence in the light of the very same original evidence. The assumption that one's peer shares one's total evidence at the point at which the disagreement manifests itself is built into our characterization of peer disagreements (see section 1). Accordingly, cases in which deletion of part of one's evidence is called for are in all likelihood not cases of peer disagreement in the intended sense. See also footnote 18 above.

[36] See also the literature of the so-called Judy Benjamin problem as originally introduced by van Fraassen (1981) and further discussed in, for instance, van Fraassen et al. (1986), Bradley (2005) and Douven and Romejin (2011). Since the Judy Benjamin problem calls for a revision of conditional probabilities, it bears interesting structural similarities to cases of peer disagreement, on our conception of them.

from our total present epistemic state.[37] One could take a different approach, and we do not want to suggest that this is the only way of implementing our proposal, but it strikes us as a particularly natural one.

The constraint we are interested in takes the form $P(H|E) = x$, for fixed H, E and x. Formally, we can identify a constraint with the set of probability functions defined on our original algebra which satisfy the constraint. In general, there will be many probability functions satisfying a given constraint. In particular, this seems to be true of the type of constraint we are presently concerned with: many initial credence distributions will be such that $P(H|E) = x$, provided that this requirement is not probabilistically incoherent, for instance because E actually implies that H is false. So, the constraint alone will not determine a unique probability function. We therefore need a way of selecting from the candidate revisions. A natural idea would be to take a probability function which departs no more from our original function than it needs to.

Skyrms (1987) suggests to describe minimal belief revisions in terms of a similarity gauge $\sigma(P, P')$, which maps two probability functions to a non-negative real number representing how much P' departs from P. An immediate requirement – *centering* - on σ would be that $\sigma(P, P') = 0$ iff $P = P'$. A given probability function does not depart from itself and every other probability function departs from it at least a little. A more complex requirement Skyrms considers is for σ to be *convex*, i.e. that it underwrites the following inequality: $\sigma(P, \lambda \cdot Q + (1-\lambda) \cdot Q') \leq \lambda \cdot \sigma(P, Q) + (1-\lambda) \cdot \sigma(P, Q')$, with λ out of the unit interval. Convexity essentially means that mixing two probability functions results in a probability function which is at least as similar to P than the least similar of the functions Q and Q'. In other words, it implies that mixing does not create out of a very similar function and a moderately similar function an unsimilar function. This appears to be plausible: mixing makes two probabilty functions to be more alike and thus makes the least similar function more similar by moving it towards a more similar one. Joyce (1999: 186f.) considers a further condition which generalizes the thought that if Q's values are uniformly closer to P's than the values of another function R, than Q should be more similar to P than R. The idea would be that similarity between probability functions is solely a matter of the distance between the values they assign.

In this setting, it is guaranteed that we will always find a unique probability function most similar to our original function out of the functions satisfying the constraint that a certain conditional probability assumes a particular value (Skyrms 1987). Moreover, the revision process from P to P' has some basic desirable features. If P satisfies already the constraint in question, then P' will be identical to P, for P is always most similar to itself (by centering). We also have the important fact that P' actually satisfies the constraint in question, for P' is chosen out of the probability functions doing so.[38]

Now, if P_{old} is our initial probability distribution and we want our new initial probability function P_{new} to be such that $P_{new}(A|E) = x$, then we can set P_{new} to be the most similar function to P_{old} satisfying this constraint. Our new belief state would then be represented by the pair (P_{new}, E), our new initial belief distribution combined with our old evidence. As a result, our new credence in a proposition H would then be $P_{new}(A|E)$. If the constraint specified this conditional

[37] This is the type of *Pluralistic Bayesianism* defended by Brössel and Eder (2014), following Levi (1980). The update mechanism proposed below diverges from theirs, however. For lack of space, a detailed comparison of both views has to wait for another occasion.

[38] The revision procedure also has a third, more complex property familiar from the semantics of conditionals. See Joyce (1999: 184).

probability to be as the Equal Weight View would have it, i.e. $P_{new}(A|E) = ½ \cdot P_{old}(A|E) + ½ \cdot P_{peer}(A|E)$, then the revision process would accord with this.

It is no accident that this revision procedure is reminiscent of the standard semantics for counterfactuals, for Skyrms has closely modelled it accordingly. In effect, probability distributions take the place of possible worlds and constraints take the place of propositions. So, instead of finding the closest possible world at which a proposition is true, we are now looking for the most similar probability function which satisfies a given constraint. The present assumptions about the similarity gauge corresponds to Stalnaker's (1968) assumption that there is always a unique world closest to the actual world. In principle, this assumption could be relaxed (by relaxing the convexity condition) to arrive at a picture which would more resemble Lewis' (1973) semantics for counterfactuals (cf. Skyrms 1987). Epistemologically, this would mean that various probability functions might be equally good candidates for being the new initial probability distribution. There would not be a single function adoption of which would be rationally mandatory. Rather, it would be rationally permissible to adopt either of the candidates.

The general question in the background is whether the revision of priors is a deterministic process or whether it leaves us with a number of possible options all of which are rationally permissible. Conditionalization is a proper updating rule that determines a unique final belief state, given the input of an initial belief state together with a piece of new information. On the other hand, it is quite plausible that the original choice of prior probabilities leaves us with more than one rationally permissible option. Although there may be only one probability function which truly represents the objective evidential relations between propositions, we have no direct epistemic access to this function. As a consequence, it does not seem irrational to adopt one out of many initial credence functions as long as certain coherence constraints are observed.[39]

In line with what we argued in section 3, our prior credences can nonetheless be understood to aim at the correct evidential relations. Being prepared to revise one's priors would then be a precondition for achieving this goal. Sensitivity to the opinions of our peers can be seen as one way of optimizing our priors, so that they have a better chance to get close to the true evidential relations in the long run.

Given this situation, it is not a far-fetched conjecture that the revision of our prior probabilities might be structurally similar to the original choice of our prior probabilities in displaying a similar kind of indeterminacy. As we have seen, it is not a mandatory assumption, though, for the model can also validate the thought that the revision procedure should be determinate. In any case, the kind of indeterminacy will be limited to those probability functions which are equally good candidates for *minimally adjusting* our old probability function. Not anything will go.

The similarity relation on probability functions presupposed in this approach would clearly require more discussion, which lies, however, beyond the scope of the present paper. Let us only present one problem case which can be taken to suggest that closeness between probability functions might not be solely a matter of the distance in values assigned to individual propositions. It seems there can be cases where the way we weigh the evidence with respect to one subject matter has implications for how we should weigh it with respect to others. For instance, a weather forecaster who moves her credence in tomorrow's weather somewhat towards that of her peer might need to change her credences also in situations which are – in a

[39] A plurality of permissible initial credence functions is often thought to be less problematic in the light of theorems that show that two probability functions tend to converge the more evidence is acquired, as long as they do not differ with respect to the propositions to which they assign probability 0 (see e.g. Gaifman and Snir 1982).

sense to be explained – structurally analogous. Suppose her credence about tomorrow's weather was based on an inductive inference from a given inductive base with a particular structure. Lowering her credence in this situation should come with the readiness to lower her credence in other situations where an analogous inductive inference was drawn from an inductive base with the same structure. Thus, when we re-evaluate an inductive inference, we are often required to think differently about what counts as a representative sample or about what counts as a statistically significant set of data. But once we change the way we draw inductive inferences, this has implications for all relevantly similar cases. If this is so, then a probability function which changes more values might actually be closer than a function which leaves more of our old credences intact. By changing more, we would achieve a more coherent set of beliefs.

In one respect, our description is still incomplete. If our new belief state were completely described by $P_{new}(\_|E)$, we would, in general, fail to believe that we disagree with our peer. The point is that our initial evidence E, which we possess before the disagreement becomes manifest, usually does not make it likely that the disagreement will occur. The situation is unlikely to change once we have adjusted our credence in the proposition we disagree about. Although we may now have a lower or higher credence in this proposition, the evidence will still not make it likely that the disagreement has occurred. Thus, we may have adjusted our credence concerning the possibility of a tsunami hitting the Malaysian coast given our original evidence, but our model does not suggest that our credence in a peer disagreeing with us about the matter has changed in any way. With respect to this fact, our evidence will in general be as strong, or as weak, as before.

The problem is not a deep one, though. What we have suggested so far is that this piece of information should make us rethink the conditional probability we originally assigned to A given E. But of course, we also need to represent the fact that we do this by learning about the proposition D. In addition to challenging our prior probabilities, peer disagreements also provide purely factual information about what one's peer says or thinks about the target proposition. As far as we can see, there is no reason to suppose that this information cannot be dealt with in the standard way by conditionalizing on it. Hence, for any proposition A, our ultimate belief state should be represented by something like $P_{new}(A|E \,\&\, D)$, rather than merely by $P_{new}(A|E)$. The evidential import of peer disagreement is thus twofold: it recommends a revision of our priors but it also provides a new piece of ordinary evidence, viz. evidence about the psychology of our peer.

8. Features of a revision of initial probabilities

On the present picture, a belief state consists of two components: a body of evidence and a probability function weighing the evidence by conditionalizing on it. Rationality does not demand to hold fixed either of these two components. It is clear that the evidence set is allowed to grow – this is the paradigmatic instance of learning. Such additions to our evidence can be modelled in terms of conditionalization. Though this lies outside the scope of our paper as it has little to do with peer disagreement, recall that the evidence set may also shrink at certain points of our epistemic life (as testified by our tendence to forget) and there may be mixed cases coming from learning something which is incompatible with our current evidence set. (The contrast between these two kinds of cases is usually drawn in terms of mere contractions and proper revisions). More importantly, we have argued that peer disagreement calls for a revision

of our initial propabilities. If this is so, the second component of a belief state can rationally be changed, too. Now, we have sketched in the previous section how one might model the revision of prior probabilities. In this section, we discuss which features this revision procedure has. Given that we are pluralists about permissible update procedures, a particularly pressing question is how a revision of prior probabilities is going to interact with conditionalization.

Fitelson and Jehle (2009) study various update rules intended to subserve the Equal Weight View. The authors consider a number of *prima facie* reasonable constraints derived from the literature on the aggregation problem, which are, however, jointly inconsistent. Given that we propose an additional update rule besides conditionalization, their discussion and the problems they raise are eminently relevant to our account. So, we will follow their approach here and go through a set of initially plausible constraints. This should help us to give a clearer picture of the update procedure we have in mind and will also point to certain issues where further research might be necessary.

A natural constraint Fitelson and Jehle discuss is what they call *probabilism*. This simply means that the end result of the update procedure should again be a probability function. This constraint is obviously satisfied on our account because only probability functions are in the domain of the similarity relation determining the new belief state. Another feature of our account is that it satisfies the constraint they call *straight averaging*: the new credence in the target proposition is – in the domain of the proper application of the Equal Weight View – the average value of the two credences initially assigned by the two peers. As a consequence, the peers will be in exact agreement after updating on the target proposition. This also means that on our account, the new credence they assign is a function of the old credences – which is the constraint Fitelson and Jehle call *irrelevance of alternatives*.

These are the features which are easily identified. A harder question is how a revision of prior probabilities interacts with the rule of conditionalization. A natural hope is that it will in some sense 'respect' conditionalization. Fitelson and Jehle interpret this as requiring that the two update rules should *commute* with each other. If we first learn A and update on it by conditionalization and then learn B and update by revising our priors, the end result should be the same if instead we first had learned B and updated on it by revising our priors and only then had conditionalized on A. Fitelson and Jehle argue that if the Equal Weight View involves straight averaging, then if it is construed as a separate update rule, it will not commute with conditionalization. The same argument is given by Wilson (2010). Sometimes the order in which we learn something will make a difference to the final result.

It is not clear whether failures of commutativity are to be expected on the present account. The important thing to note is that the identity of the information provided by peer disagreements changes as the evidence with respect to which we disagree increases. Peer disagreements are always relative to the shared body of evidence with respect to which the peers involved arrive at their conflicting evaluations of the target hypothesis H. So, initially, our disagreement with *z* may be captured by D[*z*, *x*, H, E], representing that *z* has credence *x* in H given E, where *x* differs from our own credence in H given E, while later on, after our evidence has increased, our disagreement with *z* may instead be captured by D[*z*, *x*, H, E^+], where *x* differs from our own credence in H given E^+. But these are clearly different disagreements, and so we learn different things on these two occasions: on the former occasion, doubt is cast on our credence in H condtional on E, whereas on the latter, what is questioned is our credence in H conditional on E^+. Accordingly, suppose that we start off with total evidence E. Subsequently we learn a new piece of ordinary information F and integrate it *via* conditionalization. Our body of total evidence will

now be E^+, say, and so be different from E. Now, we learn about a peer disagreement with *z* and readjust our conditional credences given E^+. Contrast this with a case in which we first learn about a peer disagreement with *z* and then learn about F. If our credences are different, this may look like a failure of commutativity. But this will only be superficially so. In the first case, we first learn F and then D[*z*, *x*, H, E^+], viz. that our peer *z* disagrees with us about the evidential import of E^+. In the second case, we first learn D[*z*, *x*, H, E], viz. that *z* disagrees with us about the evidential import of E, and then learn F. As pointed out above, what we learn at the first step in the second case is different from what we learn at the second step in the first case. Hence, the failure of commutativity is only apparent.

As we currently see it, the standard formulation of commutativity does not really apply to our case because it is stated by insufficiently specifying what we learn in the context of peer disagreements. As we have tried to explain, the disagreement is implicitly relative to the body of evidence we share. But given that the body of evidence grows when we update by conditionalization, a later disagreement about the same proposition may then provide a different piece of information. Now, given that the standard formulation of commutativity seems inapplicable to our case, one may ask whether there is a way of characterizing a more sensitive version of commutativity applicable to our account. Ultimately, we have to leave this for further research, but we can point to what we think is a promising starting point. In the present framework, peer disagreements are described as disagreements relative to our total evidence. Although this is a possible limiting case, in many scenarios the disagreement may only concern a subset of our evidence. Some things may simply have nothing to do with our disagreement, neither before nor after it manifests itself. Hence, if we have a piece of information F which has no bearing on the target proposition even when conjoined with the original evidence both before and after the disagreement occurs, then we may plausibly expect that the two updating procedures commute with each other, for F is assumed not to affect the evaluation of the target proposition in any way even if it is learned before the disagreement takes place. Incorporating the possible relativity to a mere subset of our total evidence would constitute a substantive modification of the framework, which is why we have to leave the details of this for further investigation.[40]

Conclusion

All non-dogmatist views agree that peer disagreements provide genuine counterevidence to which we should, always or on occasion, respond by changing what we believe. The most principled of these views is the Equal Weight View. It would be an attractive possibility to adopt this view while still being in a position to allow for frequent assignments of probability 1 to logical and non-logical propositions alike. This would save the Equal Weight View from unwanted sceptical consequences. To make room for this possibility, we have sketched an interpretation of probability 1 in terms of full belief, on which, unlike on the more traditional

[40] Another question Fitelson and Jehle discuss is whether an update on a peer disagreement could force new disagreements on statements we previously agreed upon, either conditionally or unconditionally. In principle, this seems to be possible on our account if the antecedent agreements are due to a certain kind of coincidence in which both peers have weighed two factors differently which then by chance made them agree upon a certain proposition. After having changed their initial probability distribution, they now may also reconsider how they weighed these two factors (compare the discussion in the previous section) and so an accidental agreement may be turned into a new disagreement.

interpretation in terms of absolute or Cartesian certainty, such assignments cease to be objectionable.

However, this combination of theses is in tension with the Conditionalization Thesis, i.e. the contention that belief revision in response to peer disagreements should always proceed by conditionalization. To devise a way out of this difficulty, we have argued that updating on peer disagreement is better construed as involving a revision of our prior probabilities. There are good independent reasons for adopting such a view. The counterevidence provided by peer disagreement presents a challenge to whether we have properly gauged the probative force of our evidence, and therefore to our prior conditional probabilities. Responding to such a challenge by revising our priors seems to be the most natural reaction.

Ultimately, we arrive at the following picture of diachronic updating. The credences of an epistemic agent are a compound P(_ |E) of her total evidence E and her prior probabilities P. Conditionalization can be said to be the standard rule of how rational belief changes over time, but there are (at least) two provisos which have to be met in order for conditionalization to be applicable. The first is that the agent has not lost any evidence over time, e.g. by forgetting it. In other words, the evidence set does not always grow: sometimes it shrinks. Furthermore, there are kinds of counterevidence, of which peer disagreement provides a paradigmatic example, that should not be construed as a mere addition to our evidence. Accordingly, a second proviso is that the new evidence does not undermine, or cast doubt upon, the way the subject has gauged the probative force of her evidence, as it is reflected in her conditional probabilities. In such a case, a revision of her prior probabilities will be called for. Thus, one's prior probability function P need not remain constant throughout one's epistemic life. As we have argued, peer disagreement testifies to this possibility: peer disagreement calls for a revision of prior probabilities.

References

AUTHOR1

AUTHOR2

Bradley, R. 2005: 'Radical probabilism and Bayesian conditioning', *Philosophy of Science* 72, 342-64.

Briggs, R. 2009: 'Distorted Reflection', *The Philosophical Review* 118, 59-85.

Brössel, P. & Eder, A.-M. 2014: 'How to Resolve Doxastic Disagreement', *Synthese* 191, 2359-81.

Christensen, D. 1991: 'Clever Bookies and Coherent Beliefs', *The Philosophical Review* 100, 229-47.

Christensen, D. 2007a: 'Epistemology of Disagreement: the Good News', *The Philosophical Review* 116, 187-217.

Christensen, D. 2007b: 'Does Murphy's Law Apply in Epistemology? Self-Doubt and Rational Ideals', *Oxford Studies in Epistemology* 2, 3-31.

Christensen, D. 2010: 'Higher-Order Evidence', *Philosophy and Phenomenological Research* 81, 185-215.

Christensen, D. 2011: 'Disagreement, Question-Begging, and Epistemic Self-Criticism', *Philosophers' Imprint* 11, 1-21.

Dietrich, F. and List, C. 2010: 'The Aggregation of Propositional Attitudes: Towards a General Theory', *Oxford Studies in Epistemology* 3, 215-34.
Douven, I. and Williamson, T. 2006: 'Generalizing the Lottery Paradox', *British Journal of Philosophy of Science* 57, 755-79.
Douven, I. and Romeijn, J.-W. 2011: 'A New Resolution of the Judy Benjamin Problem', *Mind* 120, 637-70.
Elga, A. 2007: 'Reflection and Disagreement', *Noûs* 41, 478-502.
Enoch, D. 2010: 'Not Just a Truthometer: Taking Oneself Seriously (but not Too Seriously) in Cases of Peer Disagreement', *Mind* 119, 953-97.
Feldman, R. 2011: 'Reasonable Religious Disagreement', in Goldman, A. I. and Whitcomb, D. (eds.), *Social Epistemology: Essential Readings*, Oxford: Oxford University Press, 137-57.
Feldman, R. and Warfield, T. A. 2010 (eds.): *Disagreement*, Oxford: Oxford University Press.
Field, H. 1986: 'Critical Notice: Robert Stalnaker's *Inquiry*', *Philosophy of Science* 53, 425-48.
Foley, R. 2001: *Intellectual Trust in Oneself and Others*, New York: Cambridge University Press.
Gaifman, H. and Snir, M. 1982: 'Probabilities Over Rich Languages, Testing and Randomness', *The Journal of Symbolic Logic* 47, 495-548.
Gärdenfors, P. 1988: *Knowledge in Flux: Modeling the Dynamics of Epistemic States*, Cambridge/Mass.: MIT Press.
Garber, D. 1983: 'Old Evidence and Logical Omniscience in Bayesian Confirmation Theory', *Minnesota Studies in the Philosophy of Science* 10, 99-131.
Genest, C. and Zidek, J. V. (1986): 'Combining Probability Distributions: A Critique and an Annotated Bibliography', *Statistical Science* 1, 114-35.
Ginet, C. 1980: 'Knowing Less by Knowing More', *Midwest Studies in Philosophy* 5, 151-61.
Greco, D. 2013: 'Probability and Prodigality', *Oxford Studies in Epistemology* 4, 82-107.
Greenough, P. and Pritchard, D. 2009 (eds.): *Williamson on Knowledge*, Oxford: Oxford University Press.
Hacking, I. 1967: 'Slightly More Realistic Personal Probability', *Philosophy of Science* 34, 311-25.
Hájek, A. 2003: 'What Conditional Probability Could Not Be', *Synthese* 137, 273-323.
Hájek, A. 2012: 'Staying Regular?', ms.
Harman, G. 1973: *Thought*, Princeton: Princeton University Press.
Jeffrey, R.C. 1965: *The Logic of Decision*, Chicago: The University of Chicago Press.
Jehle, D. and Fitelson, B. 2009: 'What is the "Equal Weight View"?', *Episteme* 6, 280-93.
Joyce, J. 1999: *The Foundations of Causal Decision Theory*, Cambridge: Cambridge University Press.
Kaplan, M. 2009: 'Williamson's Casual Approach to Probabilism', in Greenough and Pritchard (2009), 122-39.
Kelly, T. 2005: 'The Epistemic Significance of Disagreement', *Oxford Studies in Epistemology* 1, 167-96.
Kelly, T. 2010: 'Peer Disagreement and Higher Order Evidence', in Feldman and Warfield (2010), 111-74.
Kornblith, H. 2010: 'Belief in the face of Controversy', in Feldman and Warfield (2010), 29-52.
Kripke, S. A. 2011: 'On Two Paradoxes of Knowledge', in Kripke, S. A., *Philosophical Troubles: Collected Papers*, *Vol. 1*, Oxford: Oxford University Press, 27-51.

Lasonen-Aarnio, M. 2014: 'Higher-Order Evidence and the Limits of Defeat', *Philosophy and Phenomenological Research* 88, 314-45.
Leitgeb, H. 2014: 'The Stability Theory of Belief', *The Philosophical Review* 132, 131-71.
Levi, I. 1980: *The Enterprise of Knowledge: An Essay on Knowledge, Credal Probability, and Chance*, Cambridge/Mass.: MIT Press.
Lewis, D. 1996: 'Elusive Knowledge', *Australasian Journal of Philosophy* 74, 549-67.
Quine, W. V. O. 1951: 'Two Dogmas of Empiricism', *The Philosophical Review* 60, 20-43.
Savage, L. J. 1967: 'Difficulties in the Theory of Personal Probability', *Philosophy of Science* 34, 305-10.
Shogenji, T. 2007: 'A Conundrum in Bayesian Epistemology of Disagreement', ms.
Skyrms, B. 1987: 'Updating, Supposing, and Maxent', *Theory and Decision* 22, 225-46.
Sorensen, R. 1988: 'Dogmatism, Junk Knowledge, and Conditionals', *The Philosophical Quarterly* 38, 433-54.
Spohn, W. 2012: *The Laws of Belief: Ranking Theory and its Philosophical Implications*, Oxford: Oxford University Press.
Stalnaker, R. 1984: *Inquiry*, Cambridge/Mass.: MIT Press.
Talbott, W. J. 1991: 'Two Principles of Bayesian Epistemology', *Philosophical Studies* 62, 135-50.
Titelbaum, M. G. 2013: *Quitting Certainties: A Bayesian Framework Modeling Degrees of Belief*, Oxford: Oxford University Press
Van Fraassen, B. C. 1981: 'A Problem for Relative Information Minimizers in Probability Kinematics', *British Journal for the Philosophy of Science* 32, 375-79.
Van Fraassen, B. C. 1984: 'Belief and the Will', *The Journal of Philosophy* 81, 235-56.
Van Fraassen, B. C., Hughes, R. I. G. and Harman, G. 1986: 'A Problem for Information Minimizers, Continued', *The British Journal for the Philosophy of Science* 37, 453-63.
Van Fraassen, B. C. 1995: 'Belief and the Problem of Ulysses and the Sirens', *Philosophical Studies* 77, 7-37.
Weatherson, B. unpublished: 'Do Judgments Screen Evidence?', ms.
Weisberg, J. 2005: 'Conditionalization, Reflection, and Self-Knowledge', *Philosophical Studies* 135, 179-97.
Williamson, T. 1997: 'Knowledge as Evidence', *Mind* 106, 717-41.
Williamson, T. 1998: 'Conditionalizing on Knowledge', *British Journal for the Philosophy of Science* 49, 89-121.
Williamson, T. 2000: *Knowledge and Its Limits*, Oxford: Oxford University Press.
Williamson, T. 2005: 'Replies to Commentators', *Philosophy and Phenomenological Research* 70, 468-91.
Williamson, T. 2009a: 'Probability and Danger', *The Amherst Lecture in Philosophy* 4, 1-35. <http://www.amherstlecture.org/williamson2009/>.
Wilson, A. 2010: 'Disagreement, Equal Weight and Commutativity', *Philosophical Studies* 149, 321-26.