

What is (Dis)Agreement?

Darrell P. Rowbottom

darrellrowbottom@ln.edu.hk

When do we agree? The answer might once have seemed simple and obvious; we agree that p when we each believe that p . But from a formal epistemological perspective, where degrees of belief are more fundamental than beliefs, this answer is unsatisfactory. On the one hand, there is reason to suppose that it is false; degrees of belief about p might differ when beliefs simpliciter on p do not. On the other hand, even if it is true, it is too vague; for what it is to believe simpliciter ought to be explained in terms of degrees of belief.

This paper presents several possible notions of agreement, and corresponding notions of disagreement. It indicates how the findings are fruitful for the epistemology of disagreement, with special reference to the notion of epistemic peerhood.

1. Introduction – The Simple Account of Agreement

When do we agree concerning some proposition, p ? On the face of it, the answer is simple; we agree that p if and only if we each believe that p . This says nothing about the nature of belief, but this is a virtue. For example, those who take beliefs to be dispositions to act in particular ways, such as Price (1969), Marcus (1990) and Schwitzgebel (2002, 2010), may understand agreement as being more closely linked

to action than those who take beliefs to be mental representations, such as Fodor (1975) and Cummins (1996).¹

In contemporary formal epistemology, however, degree of belief replaces belief as the central notion. In philosophy of probability, the concept of a degree of belief is now relatively old; notably, it appears in the work of Keynes (1921), Ramsey (1926), and De Finetti (1937). However, using degrees of belief to tackle general epistemological questions is relatively new. The explanation is partly that probability, and the philosophy thereof, was once considered to be the domain of the philosopher of science. (And even within that domain, the importance of degrees of belief, as distinct from logical relations, was long disputed. Both Carnap and Popper, for example, defended logical views of probability.²) Hence, degrees of belief were typically employed only in models of scientific confirmation and reasoning (within philosophy, as opposed to the social sciences).

The reticence to use degrees of belief in broader contexts may also be partly explained by the ongoing controversy surrounding *their* nature. That is to say, there is a ‘nature of degree of belief’ debate that parallels the ‘nature of belief’ debate mentioned above. Are degrees of belief to be understood as mental representations, as dispositions to bet in particular ways, or as something else altogether? The lack of

¹ Some views on the nature of beliefs nevertheless have interesting consequences for views on agreement. For example, Schwitzgebel’s (2002, 2010) notion of ‘in-between believing’ suggests the possibility of ‘in-between agreement’. See footnote 9.

² Logical views can, but need not, involve degrees of belief. Fundamentally, they ground probabilities on logical relations between propositions. For more on this, see Rowbottom (2008) and Rowbottom (2015, pp. 24–26).

perceived progress in the debate is illustrated by the recent thesis of Eriksson and Hájek (2007), namely that ‘degrees of belief’ are basic and impervious to analysis.

But let’s take it as uncontroversial that epistemology has something to learn from formal epistemology (without presuming anything so rash as that the latter has any form of priority over the former, or that the converse is not true). And let us start with a rough and ready notion of a ‘degree of belief’, as a *degree of confidence*. A simple example captures the essence of the idea. I am more confident that I am writing a paper on agreement than I am that it will be published in a leading journal of philosophy. But I nevertheless believe each proposition to be true.

Using degrees of confidence is sufficient to provide a coarse-grained map of the terrain, and indicate that much of the land is fertile. If this makes you feel uneasy, consider that informal epistemology normally proceeds without presuming an answer to, or even considering, ‘What is belief?’ This question is understood to fall within the philosophy of mind. Hence, it would be unwise to demand that the formal epistemologist to begin with a complete answer to ‘What is degree of belief?’

So we have a plan. We will examine how we may construe agreement in terms of degrees of confidence.³ But what is the point of the exercise? Intrinsic interest aside, understanding agreement is key to understanding disagreement. (Although, as we will

³ Only agreement and disagreement between individuals will be discussed in this paper. However, several of the results plausibly hold for group level agreement and disagreement. For instance, the D-functions discussed might be construed to reflect group degrees of belief. On group degrees of belief and probabilities, see Gillies (1991), Gillies (2000, ch. 8), Rowbottom (2013) and Rowbottom (2015, ch. 6).

see, disagreement should not be equated merely with the absence of agreement.) And the epistemology of disagreement is an area of burgeoning interest, as evidenced by a number of recent edited works, e.g. Christensen (2009), Feldman and Warfield (2010), and Christensen and Lackey (2013). Participants in this debate typically operate with an intuitive understanding of disagreement. As we will see, however, analyzing disagreement proves productive.

Only a brief warning remains, before we begin. We will not work with the idea that there is a single correct account of agreement; and we will assume neither that the folk notion of agreement, nor that our intuitions about agreement, are of special significance. Rather, we will be exploring, and refining, several alternative understandings of agreement.

2. Perfect Agreement and Close Agreement

Imagine two people, A and B, have the same evidence, e , concerning whether a hypothesis, h , is true. (We will take evidence to be propositional in character, for the time being, although this is an issue that we will return to in the final section.) More precisely, imagine that each has an active degree of confidence involving these propositions, $D_A(h|e)$ and $D_B(h|e)$. Think of $D_S(h|e)$ by analogy with $P_S(h|e)$ on a subjective theory of probability; it represents how confident a subject S is about h , assuming e is true. $D_S(h|e)$ is *active* for S when e represents S's current evidence (or

background assumptions).⁴ S will also have *inactive* conditional degrees of confidence concerning h , such as $D_S(h|e^*)$, where e^* is different from e . $D_S(h|e^*)$ reflects S's degree of confidence in h conditional on S's evidence being e^* .

The motivation for using D-functions, rather than P-functions, is that we can discuss degrees of confidence without presuming that they satisfy the axioms of probability. And it appears they don't, even if they should, because some people do allow themselves to be Dutch Booked.⁵ But irrationality does not prevent agreement; two irrational individuals may agree on h , just as an irrational individual may agree with a rational individual on h .

The ground is prepared. We may begin by considering a case where agreement is clear cut. Let's call this perfect agreement:

⁴ The concept of evidence in operation is subjective, rather than objective; so keep 'background information' in mind. For more on the different notions, see Achinstein (2001), Rowbottom (2014), and Williamson (2015).

⁵ This matter isn't as simple as it may first appear. More accurately, we should say that people sometimes select betting quotients that are not coherent (or make bets such that they lose money whatever happens). In this event, their betting quotients violate the axioms of probability. But only on the assumption that the betting quotients have the same (or appropriately similar) values as degrees of confidence may we conclude that the latter do not satisfy the axioms of probability. See Gillies (2000, pp. 55–58) and Rowbottom (2015, pp. 38–52).

Perfect Agreement on h at Time t: $D_A(h|e)=D_B(h|e)$, where e is A's evidence concerning h at t and e is B's evidence concerning h at t .⁶

It is unnecessary, however, to insist that agreement (between A and B on h) is only present, at any time t , under such circumstances. Rather, it appears that phenomena worthy of the name 'agreement' can persist even in the absence of such perfect agreement.

First, let's allow the degrees of confidence of A and B to vary, in a hypothetical class of scenarios, while keeping everything else fixed. We might initially consider the following:

Close Agreement on h at Time t: $D_A(h|e)\approx D_B(h|e)$, where e is A's evidence concerning h at t and e is B's evidence concerning h at t .

In a broad range of situations, close agreement on h is indistinguishable from perfect agreement on h . Provided neither cites anything other than their evidence concerning h , an earnest discussion between A and B on the topic of h is unlikely to result in any changes to A's and B's active degrees of conviction about h . (That is, assuming A does not treat evidence that B agrees on h as evidence that h , as he might do if she

⁶ Some minor refinements might be advisable; for example, one might specify that agreement does not obtain, when evidence is shared, in the event that said evidence is internally inconsistent. For present purposes, let's handle this by stipulating that evidence may not have this feature. Other similar minor adjustments are highly unlikely to bear on what follows.

thinks B is an epistemic peer, and vice versa.⁷) Such a discussion would likely result in each party thinking that the other had reasonable, indeed correct, views about h .

Nonetheless, there are situations where the presence of only close agreement, rather than perfect agreement, is noticeable; if A and B seek to agree on a fair betting quotient on h in a high-stakes situation (for both of them), then a minor dispute between them might ensue.

3. Context and Agreement as a Threshold Difference

It is now natural to wonder how much difference there must be between $D_A(h|e)$ and $D_B(h|e)$ in order for A and B to disagree on h . Consider the magnitude $D_A(h|e) - D_B(h|e)$, i.e. $|D_A(h|e) - D_B(h|e)|$. Call this $D_{A-B}(h|e)$. What range of values may this function take when A agrees with B about h ?

To ask this question without qualification is to make two questionable assumptions, namely that disagreement is: (a) all-or-nothing, and (b) contextually invariant. Let's consider these in turn.

Keep considering a hypothetical situation where *close agreement* holds, but where A and B are in dispute about how to bet on h because the stakes are high. On the one hand, we might say that A and B agree on h but disagree about how to bet on h . On the other hand, we might say that A and B disagree about how to bet on h because

⁷ See Kelly (2005) and Elga (2007) on the wisdom of conciliation in such circumstances. For present purposes, note merely that A might *treat* B's agreement that h as evidence that h even if she should not.

they disagree about h . And it is far from obvious which path is preferable, let alone whether there is a fact of the matter about which is correct. It may therefore be better to operate directly with the notion of difference in confidence (which one could refer to as ‘degree of disagreement’), i.e. $D_{A-B}(h|e)$, in cases where background evidence is shared.

One might couple this with thinking that agreement is (typically) context relative, or at least that agreement ascriptions are (typically) context relative. In other words, the difference in confidence required for (attributing) agreement may vary according to the situation in which the actors find themselves. For example, if $D_{A-B}(h|e) \approx 0$, i.e. if so-called ‘close agreement’ holds, then A and B might be said to agree unless they find themselves in an exceptional situation, where being right about h is of extreme importance.

Remain unconvinced? The alternative is that there is some threshold value, r , such that A and B agree on h when $D_{A-B}(h|e) < r$. However, the implausibility of this can be illustrated by a consideration of limit cases. First, imagine that A is neutral about whether h or not- h , whereas B is slightly more confident that h than not- h ; $D_{A-B}(h|e) < r$, $D_A(h|e) = D_A(\sim h|e) = 0.5$, and $D_B(h|e) > D_B(\sim h|e)$. If forced to bet for or against h , A might flip a fair coin to decide. In the same situation, however, B might resist flipping a fair coin. So if A and B were forced to bet jointly (with shared funds), we would expect them to be more likely to agree on betting for h than against h , *ceteris paribus*.

Second, consider a case where A is certain that h but B is only strongly convinced that h ; $D_{A-B}(h|e) < r$, $D_A(h|e) = 1$, and $D_B(h|e) < 1$. And imagine that being wrong about h would have dire consequences. Let's say A and B are morally upright scientists, working together to prevent the extinction of the human race. They have identified two possible courses of action to deal with a virulent virus, which has infected all humans. They are certain that the first course of action will result either in all infected individuals being cured if h is true, or all infected individuals dying if h is false. They are also certain that the second course of action will result in between 70% to 90% of all infected individuals being cured (and the remainder dying). Will A and B advocate the same course of action, even assuming that they share the same values and desires?⁸ Not necessarily. A will recommend the first course of action. But B may prefer the second course of action, due purely to her sliver of doubt that h is wrong.⁹

To resist the conclusion that agreement (or ascriptions thereof) should be understood as contextual, one might instead deny that 'close agreement' ever counts as genuine agreement (but accept that 'perfect agreement', for example, does). This does violence to the folk (or everyday) notion of agreement, however, since we normally

⁸ This may be made more precise, if we think in terms of decision or game theory; we may say that the utilities for each option are the same for each scientist, but that the expected utilities for each diverge in virtue of their differing degrees of confidence.

⁹ The prior examples also provide a basis for challenging the theory that agreeing that p is simply sharing the belief that p , on some accounts of the nature of belief. Consider the dispositionalism advocated by Schwitzgebel (2002: 253), according to which: 'To believe that P ... is nothing more than to match to an appropriate degree and in appropriate respects the dispositional stereotype for believing that P.' A and B may have near identical dispositional profiles, each matching the dispositional stereotype for believing h in appropriate respects. Yet it is not obvious that they agree on h .

ascribe agreement whenever ‘close agreement’ holds. And the folk notion is useful for predictive purposes, at the bare minimum, in everyday contexts.

Hence, one might proceed with an explicitly contextual definition such as the following:

Threshold Agreement on h at Time t : $|D_A(h|e) - D_B(h|e)| < r$, where e is A’s evidence concerning h at t , e is B’s evidence concerning h at t , and the value of r is set by the context in which A and B are situated.

As we will shortly see, such a definition is partial on some reasonable conceptions of agreement – and presents only a sufficient, but not a necessary, condition for agreement – where agreement is possible even when parties (such as A and B) have different background information (‘evidence’). Beforehand, however, we should consider the possibility that agreement can be thought of in a diachronic sense.

4. Diachronic Agreement

The previous conceptions of agreement are all synchronic, although diachronic conceptions are also possible. For example, A and B might have identical *conditional* degrees of belief concerning h , such that they will be in perfect agreement on h whenever they share the same evidence. More precisely, we can state this condition as follows.

Shared Evidence Perfect Agreement on h: $\{D_A(h|e_1)=D_B(h|e_1), D_A(h|e_2)=D_B(h|e_2), \dots, D_A(h|e_{n-1})=D_B(h|e_{n-1}), D_A(h|e_n)=D_B(h|e_n)\}$, where $\{e_1, e_2, \dots, e_{n-1}, e_n\}$ is the set of all possible states of background information for A and B.

It is easy to see how one might generate other similar conceptions. For example, a weaker version of the prior condition, where A and B will only be in close agreement whenever they share the same evidence, can be outlined as follows.

Shared Evidence Close Agreement on h: $\{D_A(h|e_1)\approx D_B(h|e_1), D_A(h|e_2)\approx D_B(h|e_2), \dots, D_A(h|e_{n-1})\approx D_B(h|e_{n-1}), D_A(h|e_n)\approx D_B(h|e_n)\}$, where $\{e_1, e_2, \dots, e_{n-1}, e_n\}$ is the set of all possible states of background information for A and B.

These may be conceptualized as situations where synchronic agreement has a high degree of stability (although higher degrees are possible, e.g. in the event that the parties' degrees of confidence are insensitive to changes in evidence). However, there is often a sense in which 'agreement' on an issue implies more than synchronic perfect or close agreement, in so far as it implies that bringing the parties to disagree on that issue (synchronically) would not be an easy matter. Perhaps the difference between the locutions 'agreement that' and 'agreement on' is significant here. That is, in so far as what follows 'agreement on' need not be a specific proposition, but can instead be a label for a cluster of propositions with a vague boundary. 'A and B agree on foreign policy' and 'A and B agree on political ideology', for example, suggest that a considerable measure of stability in close agreement obtains.

More limited conceptions of diachronic agreement may be employed, to capture such cases; for example, it may be specified that only specific changes (or classes of change) in shared evidence would result in the lack of (perfect or close) agreement. And by comparing the changes, or classes of changes, rankings of stability may sometimes be possible.

Diachronic extensions of many of the following synchronic conceptions are also possible. These will not be presented, for simplicity's sake.

5. Superficial Disagreement

Consider now a specific kind of case where two people disagree on whether a hypothesis is true – in so far as one asserts that it's true, whereas another asserts that it's false – although their relevant *conditional* degrees of belief are identical or highly similar. More precisely, consider that the following two conditions may hold despite (unconditional) disagreement:

Symmetrical Conditional Perfect Agreement on h at Time t : $D_A(h|e)=D_B(h|e)$
and $D_A(h|e^*)=D_B(h|e^*)$, where e is A's evidence concerning h at t and e^* is B's evidence concerning h at t , and e is not identical to e^* .

Symmetrical Conditional Close Agreement on h at Time t : $D_A(h|e)\approx D_B(h|e)$
and $D_A(h|e^*)\approx D_B(h|e^*)$, where e is A's evidence concerning h at t and e^* is B's evidence concerning h at t , and e is not identical to e^* .

Here's a concrete example. Imagine two scientists, Alex (A) and Belinda (B), are discussing whether the Higgs boson has been detected. The evidence of A, e , is contained within the evidence of B, e^* . However, B has additional evidence of which A is unaware, from an experiment she has just taken part in; in summary, e is a proper subset of e^* . Here's a way their discussion might unfold. Belinda mentions the Higgs boson. Alex expresses the view that caution about its existence is still appropriate. In response, B expresses surprise. It emerges, in short order, that A has not 'heard the news'. So B presents the results of her recent experiment to A. As a result, A concurs that the Higgs boson exists. B then says that A was *right* to doubt the existence of the Higgs boson, before hearing about the experiment in which she had taken part.

In such situations, there is a strong sense in which disagreement about h is superficial; if A had B's evidence, then A would agree with B about h , and vice versa. Indeed, one might say that the *genuine* (or deep) disagreement concerns only the evidence for or against h ; for the resolution of *that* disagreement, in favor of either of the parties, results in their agreement on h . More precisely, the following condition obtains:

Strong Superficial Disagreement on h at Time t : $D_A(h|e) \neq D_B(h|e^*)$, but $D_A(h|e) = D_B(h|e)$ and $D_A(h|e^*) = D_B(h|e^*)$, where e is A's evidence concerning h at t and e^* is B's evidence concerning h at t .¹⁰

¹⁰ Many variants are possible, such as the following:

Weak Superficial Disagreement on h at Time t : $D_A(h|e) \neq D_B(h|e^*)$, but $D_A(h|e) \approx D_B(h|e)$ and $D_A(h|e^*) \approx D_B(h|e^*)$, when e is A's evidence concerning h at t and e^* is B's evidence concerning h at t .

Prima facie, one might think that such a situation is better described as follows: A and B agree about the extent to which they *would* individually believe in *h*, provisional on having the other's evidence, although they disagree on *h* and the evidence itself. This is incorrect, however, because it is possible for people to have false beliefs about their provisional degrees of belief, and even their active degrees of belief. Accept that there are some circumstances where degrees of belief are accurately measured, e.g. by gambling scenarios. How one finds oneself betting may deviate from how one anticipates betting, as Ramsey (1926) shows. In short, degrees of belief are not, in general, luminous. Conditional agreement, as defined, is therefore interestingly distinct from the accounts of agreement previously discussed. It need not be symmetrical, as it was in the example above. We can also recognize the existence of asymmetrical cases.¹¹

6. Superficial Agreement

Agreement may also be superficial; for example, A and B may accept the same betting odds as fair for a given event, but for different reasons. Annabel (A) and Bruce (B) might be equally confident that a particular horse will win a race, although

¹¹ For instance, we could consider the following (along with an obvious alteration for close, rather than perfect, agreement):

Asymmetrical Conditional Perfect Agreement on h at Time t: Either $D_A(h|e)=D_B(h|e)$ & $D_A(h|e^*)\neq D_B(h|e^*)$ or $D_A(h|e^*)=D_B(h|e^*)$ & $D_A(h|e)\neq D_B(h|e)$, where *e* is A's evidence concerning *h* at *t* and *e** is B's evidence concerning *h* at *t*, and *e* is not identical to *e**.

[The 'either... or' formulation denotes an *exclusive* or.]

A's confidence might stem from data about the horse's form relative to the oppositions', over recent races, whereas B's confidence might stem from insider information that the race will be fixed. In such a scenario, we may say there is:

Superficial Perfect Agreement on h at Time t: $D_A(h|e)=D_B(h|e^*)$, where e is A's evidence concerning h at t and e^* is B's evidence concerning h at t , and e is not identical to e^* .

And similarly, provided the betting odds accepted as fair are highly similar, we may say that the following holds:

Superficial Close Agreement on h at Time t: $D_A(h|e)\approx D_B(h|e^*)$, where e is A's evidence concerning h at t and e^* is B's evidence concerning h at t , and e is not identical to e^* .

It is also plausible that superficiality comes in degrees, and that by appropriately narrowing down the class of superficial agreements (either perfect or close), we may identify such degrees. Consider the following definition, which captures a proper subset of superficial perfect agreements, to see this:

Independent Perfect Agreement on h at Time t: $D_A(h|e)=D_B(h|e^*)$, where e is A's evidence concerning h at t and e^* is B's evidence concerning h at t , and e shares none of the content of e^* .

(There is an equivalent proper subset of close agreements, where ‘=’ is replaced by ‘≈’; for concision, however, we will not cover this equivalent, or similar equivalents, from this point onwards.) Now it seems that the superficiality of the agreement would be less, provided that some the content of e and e^* were to overlap (and that e and e^* were to remain consistent). Indeed, it seems reasonable to think that the superficiality of the agreement is a function of the similarity relation between e and e^* .

We will not go into the issue of how to measure similarity here. (It is a notorious problem, which plagued Popper’s account of verisimilitude. See Miller 1974 and Tichý 1974.) Suffice it to note that disagreement may be *rational* even when e and e^* differ by one proposition; a single flip of a coin where it lands on its side, for example, is sufficient to illustrate that the probability of heads or tails, on said flip, cannot have been one. (That is, provided the probability in question is not taken to be the frequency of heads or tails in the limit of an infinite collective of flips.)

7. The Results Applied: Epistemic Peerhood

We have covered a number of ways in which to make the notions of agreement and disagreement more precise, and also considered several ways by which to construe the extent of agreement or disagreement. We’ll now use these results to illuminate a key concept in the epistemology of disagreement over the past decade.

A central question in the literature on disagreement is ‘What should one do when one learns that one disagrees with an epistemic peer?’ And how this should be answered depends crucially on how ‘epistemic peer’ is understood, although, as we will see,

several authors construe it differently. Kelly (2005, p.167) reports taking the term ‘epistemic peers’ from Gutting (1982, p. 83), who specifies only that this refers to individuals with equivalent ‘intelligence, perspicacity, honesty, thoroughness, and other relevant epistemic virtues’. For Kelly (2005, p. 167), however, ‘epistemic peers’ must also be ‘equals ... with respect to their exposure to evidence and arguments which bear on the question at issue.’ More precisely, Kelly (2005, section 2.3) writes:

[T]wo individuals are epistemic peers with respect to some question if and only if they satisfy the following two conditions:

- (i) they are equals with respect to their familiarity with the evidence and arguments which bear on that question, and
- (ii) they are equals with respect to general epistemic virtues such as intelligence, thoughtfulness, and freedom from bias.

Let’s now consider these conditions in terms of the distinctions made in the previous sections. Let A (Alice) and B (Bob) be two peers in Kelly’s sense, and let the ‘question’ they are concerned with be whether a specific hypothesis, h , is true. Let e represent Alice’s evidence, and e^* represent Bob’s evidence.

Now consider the component of condition (i) relating to evidence. (We will come to arguments shortly). First, note that it is undesirably vague, in so far as what ‘familiarity with the evidence’ involves is unspecified. In fact, the problem is twofold: it’s unclear what counts as ‘evidence’, and it’s unclear, even if this is decided, what

‘familiarity’ (and hence equality in measures thereof) involves. So let’s explore some options.

Let’s assume that Kelly intends ‘evidence’ to be propositional (or sentential) in character, in line with the previous treatment in this paper. (If evidence is construed in a broader sense – e.g., if it may include physical objects such as guns and DNA samples – then it’s easy to see how legitimate disagreements may arise. Observation is theory-laden, and so forth.) Even then, ‘familiarity’ with a given item of evidence (for or against a hypothesis h) does not entail taking it *to be* evidence, on an objective account of evidence. For example, Alice might be familiar with the fact that there is a dark band in between primary and secondary rainbows without realizing that this is evidence in favor of the geometrical optical theories employed by Descartes and Newton. And this does not depend on her being ignorant of said theories; on the contrary, she may have studied them at secondary school.

Moreover, the use of the definite article in ‘the evidence’ seems at odds with the notion that evidence has a subjective element. And evidence *may* have such a character on several views, e.g. that S’s evidence is what S knows (Williamson 1997), and hence believes, a.k.a. $E=K$, or that S’s evidence is what S assumes to be true (Williamson 2015), a.k.a. $E=A$. Thus, Kelly presumably has in mind the idea that there is *public evidence*, or something similar but with somewhat more limited scope, in line with Williamson (1997, p.717):

Science depends on public evidence, which is neither the union nor the intersection of the evidence of each scientist. “It is known in S” and “We

know” (which is not synonymous with “Some/many/most/all of us know”) express the corresponding kind of knowledge.¹²

Thus, we may understand ‘equals with respect to their familiarity with the evidence’ as ‘equals with respect to their familiarity with the public evidence [or some surrogate]’. This resolves the prior concerns, if we accept that what *counts* as public evidence pertinent to *h* is clear to Alice and Bob, e.g. in virtue of the source(s). We are now left to wonder exactly how to cash out ‘equals with respect to... familiarity’. And this will prove to be a context dependent matter, for as Kelly (2005, section 2.3) notes:

[I]nasmuch as classes of epistemic peers with respect to a given issue consist of individuals who are ‘epistemic equals’ with respect to that issue, whether two individuals count as epistemic peers will depend on how liberal the standards for epistemic peerhood are within a given context. That is, whether two individuals count as epistemic peers will depend on *how much* of a difference something must be in order to count as a *genuine* difference, according to the operative standards.

So a highly demanding standard for peerhood will require that *e* is identical to *e**, whereas less demanding standards will require lesser degrees of similarity between *e* and *e**. This is interesting, in terms of the previous analysis, in so far as for any less demanding standard, disagreement may be (either weakly or strongly) *superficial* on

¹² Such ‘public evidence’ could consist merely of statements classified as true, rather than statements that are genuinely true, on views such as $E=A$.

Kelly's account, as it stands. (And presumably the highly demanding standard will *not* be operating in most *real* cases in which one takes another to be an epistemic peer. For instance, I do not think that any other philosopher working on disagreement takes into account *precisely* the same evidence on the matter as I do.) Moreover, an individual (who knows that she is) disagreeing with an epistemic peer should take this fact into account. Small differences in evidence are enough to mask the presence of *symmetrical conditional close (or even perfect) agreement*. And the probability that such differences in evidence exist for any two peers (in a suitably liberal sense) is high. Thus it appears appropriate, at the bare minimum, to assign a non-zero probability to the fact that such conditional agreement is present. If the possibility is judged significant, indeed, then this provides reasonable grounds for failing to conciliate. A better way to proceed might be to try to tease out exactly how one's evidence differs from that of one's peer, although this will be an extremely difficult task in some cases. Think of this from Alice's perspective. On learning that she disagrees with Bob on *h*, she might, in virtue of classifying him as an epistemic peer (in a somewhat liberal sense), fail to conciliate because she takes this to likely result from some significant difference in the evidence of each, *e* and *e**, despite their shared access to the public evidence (or a proper subset thereof).

This brings me to Kelly's mention of 'arguments'. Clearly, these are propositional (or sentential) in character; and hence, we can sink these into *e* and *e**, from a formal perspective. However, familiarity with an argument does not entail accepting its conclusion. One might not accept any number of its premises. Or one might see it as a *reductio*. Or one might think that it is fallacious, or weak. And so forth. In effect, this means that Alice should bear in mind that her disagreement with Bob on *h* may be

superficial even if they are aware of precisely the same arguments said to bear on h , simply because they have different beliefs about those arguments.¹³ This reinforces the above finding about the wisdom of non-conciliation in some possible scenarios, *when peerhood is understood in Kelly's sense*. (I will use 'conciliate' loosely, as follows. To conciliate, upon learning of disagreement concerning h , is to alter one's (degrees of) belief as a result.)

So far, we have discussed only Kelly's first condition for epistemic peerhood. Let's now move on to the second, according to which equality is necessary in 'general epistemic virtues such as intelligence, thoughtfulness, and freedom from bias' (Kelly, section 2.3). What effect should we expect such equality to have? To answer this, let's assume that we're comparing two individuals with access to *identical* evidence and arguments. Let's take this condition to hold for Alice and Bob, and consider their degrees of belief concerning the members of a finite set of hypotheses, $\{h_1, \dots, h_n\}$. Should we expect them to have *close agreement* or even *threshold agreement* on each item, or at least most or many items, in that set? No. Imagine that Alice and Bob are both severely lacking, to the same extent, in general epistemic virtues. They just don't care, or think very hard, about the hypotheses in question. What's more, they have biases that they indulge. Perhaps Alice is a democrat, and Bob is a republican, and each hypothesis in the group is supported by one political party and rejected by the other. Alice just goes with whatever the democrat line is, whereas Bob goes with

¹³ What exactly it takes for an argument (to be thought) 'to bear' on a claim is also worthy of analysis. For present purposes, however, note simply that there are clear cases where two people may both think that an argument bears on h , but in a different way. For example, one may take it to confirm h , whereas another may take it to disconfirm h .

whatever the republican line is. They are equally lacking in ‘freedom from bias’, although their biases are different.

But now imagine that Bob and Alice were instead ‘epistemic equals’ in so far as they were maximally intelligent and thoughtful, totally free from bias, and so forth. Then it would be more reasonable to expect them to have close agreement or threshold agreement on some hypotheses in the set. That is, even though no such close agreement or threshold agreement is entailed unless many more assumptions – concerning the truth of evidentialism, the scope for rational degrees of belief to differ in value when evidence is fixed, and so on – are made. (It is no accident that Kelly (2010) begins with three examples where both parties are highly epistemically virtuous, and very well-informed. These examples make disagreement seem somewhat puzzling. But peerhood, as he defines it, doesn’t make disagreement between epistemic peers puzzling in general. That is, as we’ve seen, even when the standards for peerhood are not ‘liberal’.)

It should already be clear, from these examples, that it is extremely difficult to determine how to relate being ‘equals with respect to epistemic virtues’ to (probable) synchronic agreement state. And this is so even when shared evidence and arguments are assumed to be identical. The most respectable way to proceed would be to consider differences in just one virtue at a time, rather than ‘epistemic virtues’ as a bunch. But we cannot do that, of course, because Kelly is not specific about the list of relevant virtues. The possibility space under consideration is vast, and vaguely circumscribed. It would be imprudent to conclude that conciliation is in general a correct (or incorrect) strategy, on the basis of considering this space directly. It is

possible that conciliation is sometimes obligatory and sometimes impermissible; moreover, there may even be cases where it is permissible but not obligatory.

Now consider, by way of comparison, the view of Elga (2006). He acknowledges the significance of the factors used by Kelly (2005) to characterize peerhood, in so far as he recognizes that one might consider another's '*judgment*' – 'the manner in which she forms opinions on the basis of even her information' (Elga 2006, section 2) – as well as her relative state of '*information*'. However, Elga's (2006: section 10) *definition* of peerhood makes no mention of these factors whatsoever:

My use of the term 'epistemic peer' is nonstandard. On my usage, you count your friend as an epistemic peer with respect to an about-to-be judged claim if and only if you think that, conditional [sic] the two of you disagreeing about the claim, the two of you are equally likely to be mistaken.¹⁴

I take it to follow that A and B *are* epistemic peers, concerning some claim *h*, provided that A and B are equally likely to be mistaken about *h* if they disagree about *h*.

As it stands, though, this definition is curious. For example, it entails that you and I are epistemic peers concerning some claim, on which we disagree, if we're both sure

¹⁴ Elga (2006) continues: 'In defense of my use, suppose that you think that conditional on the two of you disagreeing about a claim, your friend is more likely than you to be mistaken. Then however intelligent, perspicacious, honest, thorough, well-informed, and unbiased you may think your friend is, it would seem odd to count her as an epistemic peer with respect to that claim, at least on that occasion. You think that on the supposition that there is disagreement she is more likely to get things wrong.'

to be mistaken about it. And it's highly plausible that we're both sure (for all practical purposes) to be mistaken when it comes to having accurate *degrees of belief* concerning some propositions. That is, although it may nevertheless be true that one of us is liable to have a more appropriate degree of belief than the other.

Imagine that we are each presented with the same scientific papers on global warming, and that neither of us has studied global warming seriously beforehand. However, you have studied physics extensively, and have a degree in the subject, whereas I failed science at school. We read the papers carefully. You emerge highly confident that global warming (due to human action) is a genuine phenomenon (*h*), whereas I emerge rather less confident. (So we are neither in perfect agreement nor in close agreement on *h*.) An ideal agent, however, would have a higher degree of belief than either of us on *h*, after having been exposed to the same papers; she would be extremely confident that *h*. Now add that both of us lacked (and will always lack) the capacity to be extremely confident that *h*, on the basis of reading those papers, due to our intellectual limitations. There is a significant sense in which we were equally as likely to be mistaken about *h*, under the circumstances, although it does not follow that we are, or were, epistemic peers. Your degree of belief may have been likely to be *considerably closer* to that of the (hypothetical) ideal agent than mine. And it would be wise for us to prefer your judgement on *h*, in such a situation, if we became aware of our respective backgrounds.

There are also some interesting problems with the use of probability in Elga's definition. Consider, for example, employment of a relative frequency view of probability (as is common in statistics). Then to speak of the probability of

disagreeing on a specific claim is elliptical, strictly speaking; there are probabilities only for events within collectives, rather than single events. (As von Mises (1928: 18) pithily put it: ‘FIRST THE COLLECTIVE – THEN THE PROBABILITY’.) And this raises the question of which collective Elga wishes us to consider. For example, you may have a higher probability than me of being mistaken about a particular claim, on which we will disagree, when we consider the relative frequency of your being mistaken over *all* the propositions on which we might (hypothetically) disagree. But if we consider just the cases in which we might (hypothetically) disagree about climate science – to extend the previous example – then you may have a much *lower* probability than me of being mistaken about the aforementioned claim. This ‘reference class’ problem is non-trivial, for reasons explained by Gillies (2000, pp. 119–125 & 182–183) and Rowbottom (2015, pp. 109–111). And if Hájek (2007) is correct, it cannot be avoided merely by rejecting the relative frequency interpretation of probability (which I used for illustrative purposes).

A simple way to deal with both of the aforementioned issues is to consider the extent to which diachronic forms of agreement, such as shared evidence *conditional* agreement, obtain.¹⁵ Doing so has the twin virtues of making the reference class explicit (and thus *contextualizing* the talk of ‘epistemic peers’) and making it clear to what extent the hypothetical and actual degrees of belief of ‘peers’ may differ (and thus rendering the talk of ‘epistemic peers’ more precise).

¹⁵ This is compatible with, but does not require, thinking in terms of belief-revision rules following Lam (2011). It also doesn’t require construing peerhood in terms of ‘reliability’, measured in either of the formal ways that Lam (2011) discusses.

8. Conclusion

The first half of this paper covered a variety of ways by which to make the notions of agreement and disagreement more precise. The second half illustrated that this conceptual apparatus may be put to good use in the epistemology of disagreement, and in discussions of epistemic peerhood in particular.

Acknowledgements

This paper results from an ongoing General Research Fund project on ‘Computational Social Epistemology and Scientific Method’ (#341413), funded by Hong Kong’s Research Grants Council. I’m grateful to David Christensen and Paisley Livingston for comments on an earlier version.

References

- Achinstein, P. 2001. *The Book of Evidence*. Oxford: Oxford University Press.
- Christensen, D. 2007. ‘Epistemology of Disagreement: The Good News’, *Philosophical Review* 116(2), 187–217.
- Christensen, D. (ed.) 2009. *Special Issue: The Epistemology of Disagreement. Episteme: A Journal of Social Epistemology* 6(3).
- Christensen, D. and J. Lackey (eds). 2013. *The Epistemology of Disagreement: New Essays*. Oxford: Oxford University Press.
- Cummins, R. 1996. *Representations, Targets, and Attitudes*. Cambridge, MA: MIT Press.
- De Finetti, B. 1937. ‘Foresight, its Logical Laws, its Subjective Sources’, in H. E. Kyburg and H. E. Smokler (eds) *Studies in Subjective Probability*. New York: Wiley, pp. 93–158.
- Elga, A. 2007. ‘Reflection and Disagreement’, *Noûs* 41, 478–502.

- Erikkson, L. and Hájek, A. 2007. ‘What Are Degrees of Belief?’, *Studia Logica* 86, 183–213.
- Feldman, R. and T. Warfield (eds). 2010. *Disagreement*. Oxford: Oxford University Press.
- Fodor, J. 1975. *The Language of Thought*. Cambridge, MA: Harvard University Press.
- Gillies, D. 1991. ‘Intersubjective Probabilities and Confirmation Theory’, *British Journal for the Philosophy of Science* 42, 513–533.
- Gillies, D. 2000. *Philosophical Theories of Probability*. London: Routledge.
- Gutting, G. 1982. *Religious Belief and Religious Skepticism*. Notre Dame: University of Notre Dame Press.
- Hájek, A. 2007. ‘The Reference Class Problem is Your Problem Too’, *Synthese* 156, 563–585.
- Kelly, T. 2005. ‘The Epistemic Significance of Disagreement’, in J. Hawthorne and T. Gendler (eds), *Oxford Studies in Epistemology, Volume 1*, pp. 167–196. Oxford University, Oxford.
- Kelly, T. 2010. ‘Peer Disagreement and Higher Order Evidence’, in R. Feldman and T. Warfield (eds), *Disagreement*. Oxford: Oxford University Press, pp. 111–174.
- Keynes, J. M. 1921. *A Treatise on Probability*. London: Macmillan.
- Lam, B. 2011. ‘On the Rationality of Belief-Invariance in Light of Peer Disagreement’, *Philosophical Review* 120, 207–245.
- Marcus, R. B. 1990. ‘Some Revisionary Proposals about Belief and Believing’, *Philosophy and Phenomenological Research* 50, 132–153.
- Miller, D. W. 1974. ‘Popper’s Qualitative Theory of Verisimilitude’, *British Journal for the Philosophy of Science* 25, 166–177.
- Ramsey, F. P. 1926. ‘Truth and Probability’, in H. E. Kyburg and H. E. Smokler (eds), *Studies in Subjective Probability*. New York: Wiley, pp. 63–92.
- Rowbottom, D. P. 2008. ‘On the Proximity of the Logical and “Objective Bayesian” Interpretations of Probability’, *Erkenntnis* 69, 335–349.
- Rowbottom, D. P. 2013. ‘Group Level Interpretations of Probability: New Directions’, *Pacific Philosophical Quarterly* 94, 188–203.
- Rowbottom, D. P. 2014. ‘Information Versus Knowledge in Confirmation Theory’, *Logique et Analyse* 226, 137–149.
- Rowbottom, D. P. 2015. *Probability*. Cambridge: Polity Press.

Schwitzgebel, E. 2002. ‘A Phenomenal, Dispositional Account of Belief’, *Noûs* 39, 249–275.

Schwitzgebel, E. 2008. ‘Acting Contrary to Our Professed Beliefs, or the Gulf Between Occurrent Judgement and Dispositional Belief’, *Pacific Philosophical Quarterly* 91, 531–553.

Tichý, P. 1974. ‘On Popper’s Definitions of Verisimilitude’, *British Journal for the Philosophy of Science* 25, 155–160.

von Mises, R. 1928. *Probability, Statistics and Truth*. 2nd ed. [1968.] London: Allen and Unwin.

Williamson, J. 2015. ‘Deliberation, Judgement and the Nature of Evidence’, *Economics and Philosophy* 31, 27–65.

Williamson, T. 1997. ‘Knowledge as Evidence’, *Mind* 106, 717–741.