

האוניברסיטה העברית בירושלים
THE HEBREW UNIVERSITY OF JERUSALEM

**EPISTEMIC PEERAGE, DISAGREEMENT,
AND BELIEF REVISION¹**

By

AMIR KONIGSBERG

Discussion Paper # 583 July 2011

מרכז לחקר הרציונליות

**CENTER FOR THE STUDY
OF RATIONALITY**

Feldman Building, Givat-Ram, 91904 Jerusalem, Israel
PHONE: [972]-2-6584135 FAX: [972]-2-6513681
E-MAIL: ratio@math.huji.ac.il
URL: <http://www.ratio.huji.ac.il/>

Epistemic peerage, disagreement, and belief revision¹

Amir Konigsberg

The Center for the Study of Rationality, The Hebrew University in Jerusalem

Abstract: recent debates have centred on the normative influence epistemic peerage should have on the regulation of beliefs in cases of disagreement. A dominant position in this debate is that acknowledging an epistemic peer's possession of a belief contrary to one's own ought, in itself, to lead to the revision of one's doxastic commitments. In what follows I aim to challenge and rethink the notion of peerage underlying the disagreement debate and thus reveal that the traditional view of peerage rests upon an idealized conception of similarity between disagreeing parties, and thus to show that the normative constraints derived from it are equally idealized. Constructively, I will suggest a commonsensical solution to the disagreement problem based on what I propose as a soft, more moderate conception of peerage.

Introduction

Supposedly, contrary beliefs possessed by peers serve as higher-order evidence (Kelly 2009; Christensen 2010) impinging on the validity of each of their existing belief states. Now while agreeing that acknowledging epistemic peerage may substantiate symmetric *expectations* from our peers, I aim to show that peerage does not extend to specific *instances* of disagreement. The reason being that while epistemic expectations may result from symmetric epistemic credentials such as past epistemic success and shared epistemic virtues, instances of disagreement do not ensure these

¹ Acknowledgement: an earlier draft of this paper was presented at the Center for the Study of Rationality's 2011 annual retreat in Eilat, Israel. I would like to thank the participants of the conference for their insightful comments. In particular I would like to thank Alon Harel for fruitful discussions, comments and feedback.

same credentials are applicable measures for ascertaining peerage with regard to them. Furthermore, as will be shown in the following sections, it seems that there is no way that such credentials *could* in fact be met for any particular instance. In showing this, my goal will be to soften the normative demand for belief revision in cases where there is disagreement with epistemic peers.

But this is not to say that I do not think belief revision upon discovery of disagreement is a healthy strategy; I even think that there are cases when peer disagreement should be sought, as a corrective and strategically beneficial risk-mitigating measure (Christensen 2007, p.216). In this sense I shall claim that proportional modification of belief in light of disagreement can be a productive strategy that ought to be implemented in some cases depending on one's level of conviction about the facts and the probative force of evidence.

I begin by outlining the problem of disagreement and the notion of epistemic peerage inherent to it. I go on to present two recent suggestions relating to the normative implications that such disagreements ought to have. The first will advocate that disagreement between peers demands revision of one's beliefs. The second that disagreement between peers might demand revision in some cases but not in others, and that there is therefore no encompassing normative entailment from disagreement. While in general I agree with the second of these views, I seek to show that the very idea of peerage is a notion that designates a *domain* of mutual expectation rather than ensuring identical beliefs within that domain. I will argue in favour of revising the notion of peerage while delimiting its normative implications.

Delimiting the domain of agreement

Recent debates invoking the notion of epistemic peerage and the latest surge of literature on disagreement appealing to it² seem to leave the notion itself rather under-explained. The problem of disagreement is typically formulated as stemming from a situation that demands normative intervention; a situation in which two or more agents of equal credentials, abilities, and qualifications respond differently to a shared body of evidence. The question typically debated is whether each of these agents ought to change something about his belief upon discovery of a peer's contrary belief about the same body of evidence.

While the problem thus formulated does indeed raise a question relating to the appropriate normative response to the discovery of such new information, its application to real-life cases across all domains in which disagreement can occur is questionable. Not only does the conception of peerage underlying the problem of disagreement lean towards an abstract idealization of real cases of disagreement, but even if it does in fact apply to some empirical domains, establishing peerage in these domains is not as neat and tidy as is suggested by the notions on the basis of which its idealized characterization seems to be founded. As we shall see, this may have implications for the strength of the normative implications resulting from disagreement in different domains³.

² See: (Elga 2007; Feldman 2006; Feldman 2007; Kelly 2005; Kelly 2008; Kelly 2009; Feldman & Warfield 2010; Elga 2009; Christensen 2010; Enoch).

³ Enoch (Enoch, pp.18-22) seeks to disentangle an ambiguity between three possible readings of peerage, or of 'being a peer.' Noticing this ambiguity hints that he recognizes that the notion is richer and more complex than recent accounts assume. Yet the question of what goes into established peerage across different domains and their correlated normative import is not pursued in his account.

Before continuing, let us briefly look at two recent approaches to the problem of disagreement. Both these approaches see the problem as relating to the appropriate doxastic response to discovering that one's peer holds a contrary belief. Each argues for a different normative response. The first view argues that upon discovery of disagreement, we should meet our interlocutor half way. This view is represented in the literature by Adam Elga's Equal Weight View (Elga 2007) and by David Christensen's "split the difference" position (Christensen 2007), which entails a symmetric conversational approach to disagreement (Elga 2007; Christensen 2007; Feldman 2006). The second view argues for a moderate ad hoc revision of belief based on the probative force of first and second-order evidence (Kelly 2009; Enoch). Broadly speaking, its main contention is that it isn't obvious a priori that giving equal weight to your peer's belief upon discovery of disagreement is the appropriate thing to do.

Recap: the Equal Weight View and the Total Evidence View

According to the meeting halfway view (at least in Elga's representative version of it (Elga 2007)), upon finding out that your peer disagrees you should revise your degree of confidence in your own belief so that it equals "your prior conditional probability that you would be right" (Elga 2007, p.490). Now because you are peers, this means that your confidence in your opinion ought to be no more than 0.5 of what it was before you found out about the disagreement. Where 0.5 is the probability of your being right (*qua* peer) prior to the disagreement (Enoch, p.2).

As Enoch (Enoch) rightly states, such a view gives rise to what Elga refers to as 'spinelessness.' Namely, to the implausible conclusion that one's own opinions have no weight of their own; that according to this view, upon disagreement one ought to

give up one's view much faster than one would actually be expected or willing to give it up in light of one's own relation to the evidence one has for one's belief (first-order evidence). Thus according to Elga's Equal Weight View, upon recognizing the disagreement with one's peer one ought to split the difference between one's own belief and that of one's peer, however strong one's first-order considerations may be. But as Enoch rightly asks, in such a case, now that you have discounted your previous opinion on the basis of the evidence of your peer's disagreement, "where has all the other evidence gone? (Enoch, p.16)." Where is the evidence because of which and in view of which you believed as you did prior to the discovery of disagreement? It is in response to this question that the Total Evidence View (Kelly 2009) proposes an alternative solution.

According to the total evidence view, proposed by Kelly (Kelly 2009) and endorsed in a modified version by Enoch (Enoch), in cases of peer disagreement, "what it is reasonable to believe depends on both the original, first-order evidence *as well as* on the higher-order evidence that is afforded by the fact that one's peers believe as they do" (Kelly, in: Feldman & Warfield 2010, p.142). The appeal of this view is that it doesn't just consider the beliefs which each of the peers hold. Rather it also considers the first-order (Enoch 2011, p.17) reasoning and evidence on which these beliefs are based. And it is its encompassing consideration of both the beliefs of each of the parties as well as the evidence on which these beliefs are based that it is referred to as the Total Evidence View, in that it considers so-called first and higher-order (or second-order) evidence. Yet while this view may seem more reasonable than the Equal Weight View, it is problematic on two counts: firstly, as Enoch has shown, the Total Weight View may be accused of 'double counting' (Enoch, p.18), in that what it seemingly refers to as first-order evidence is supposedly already accounted for in the

beliefs that each of the peers actually holds. So if I believe P because of E , and you believe $not-P$ because of E , each of our beliefs already includes our first-stage relation to the evidence that warrants this belief. Secondly, by favoring the first-stage evidence of one party upon the other, the Total Evidence View, proposed by Kelly, endorses what seems to be an unfounded and unlawful asymmetry favoring one agent's epistemic relation to first-order evidence to the other (Enoch, p.18). After all, since these agents are peers there ought to be symmetry in the likelihood of their evaluating evidence correctly.

Thus each of the views presented briefly above possesses some appeal, each also possesses significant problems that are highlighted in their underlying rational assumptions, some of which I believe can be untangled by attending to the notion of 'peerage' prior to looking into to the appropriate mode of revision upon discovery of disagreement.

Thus while the problem of disagreement has typically related to what an agent discovering disagreement with his peer *ought* to do with his own beliefs, I aim to discuss a related, underlying, question, about the scope and the meaning of epistemic peerage and the legitimacy of upholding it in non-idyllic, real-life cases. This is what I shall attend to in the next section.

Engaging the normative challenge posed by disagreement

There is undoubtedly a normative question of what to do when two people identical in all their abilities and prior experience disagree with one another. All the same it is questionable whether there exists a domain where this sort of symmetrical identity subsists. While we may envisage two agents differentiated only by the occasion of their present disagreement and identical in all else, following this thought through

seems to inevitably lead to its own implausibility. Since for any two agents that are completely alike in epistemically relevant composition, any doxastic deviation of one from the other in relation to the same evidence must seemingly be attributed to non-rational involuntary causes, or to whim.

Yet all the while, disagreements do arise, and they arise in rational people thought to be equal in expertise, knowledge, and proficiency, in specific domains ranging from weather forecasting (Elga 2007; Christensen 2007; Kelly 2009), economics, chess, politics (Feldman 2006), morality and justice (Rosen 2001)⁴, to name but a few. Therefore in light of the fact that peers, conceived of in this sense, do exist, it may make sense to relax the exactitude of the notions by which peerage is typically established, so as to allow for a rational conception of disagreement according to which parties can disagree while neither one must necessarily be irrational in doing so.

⁴ It is worth quoting fully a passage from Rosen (2001) expressing that disagreement between peers need not be taken as a sign of the necessary irrationality of at least one of the parties involved: “It should be obvious that reasonable people can disagree, even when confronted with a single body of evidence. When a jury or a court is divided in a difficult case, the mere fact of disagreement does not mean that someone is being unreasonable. Paleontologists disagree about what killed the dinosaurs. And while it is possible that most of the parties to this dispute are irrational, this need not be the case. To the contrary, it would appear to be a fact of epistemic life that a careful review of the evidence does not guarantee consensus, even among thoughtful and otherwise rational investigators”(Rosen 2001, pp.71-72).

Thus while peerage does indeed exist, what goes toward substantiating it differs in rigor and form depending on the domain in which such peerage is recognised⁵. This, at least is what I would like to claim. Let me therefore distinguish between *peerage* and *evidence concerning peerage*. Since while the former may serve as a description of a relation that holds between agents within or across domains while substantiating the normative implications of such relations, what goes into establishing peerage in different domains varies widely. Furthermore, it seems that the nature of the evidence establishing peerage in a domain may not only impinge on the normative entailments of disagreements between peers in that domain but may rather also relate directly to the rigor of such normative entailments.

In the most trivial sense this can be seen in that if upon disagreement with a peer one ought to revise one's existing belief to the n th degree, then the belief resulting from such a revision would be normatively sanctioned only if one's prior belief was itself normatively sanctioned. Revising a belief by meeting a contrary belief half way may be the appropriate method for revision, but the consequent belief is only appropriate if the initial belief, prior to revision, was itself appropriate. In this sense if there are different criteria establishing peerage in different domains then the justification of these criteria *qua* criteria for peerage is directly correlated with the justification of any belief revised in accordance with its own normative ruling⁶.

⁵ Enoch (Enoch n.d., p.41) alludes to the role that the "factors that differ from one case of peer disagreement to another" in establishing the weight of a peer's disagreement in revising our belief. I read Enoch's point as alluding to the point I am expanding on here.

⁶ This point is taken from Enoch's call for a revision of Elga's version of the Equal Weight View (Enoch, p.20). Enoch argues that Elga's account of the Equal Weight View is ambiguous in that it leaves open the justification of peerage, which it assumes.

Evidence for peerage

When we say that two people are peers *because of* their shared or common credentials, we may, in most domains in which it is plausible that peerage is established, be referring not to equal epistemic or cognitive capacities nor to their equal past experience. Rather, in most cases we will be conferring that they have equal *epistemically relevant* capacities or equally *relevant* past experience. Where ‘equally relevant’ past experience can here be conceived not in terms of the identity of the experiences of different agents but in terms of their having had like success on a number of relevant measures, such as equal scoring (e.g., six out of ten) on the relevant tests, or an equal number of prior attempts at some particular type of quiz (e.g., four attempts). In cases of the latter kind the factors according to which peerage is established are more relaxed than those assumed when the notion of peerage is idealized. In idealized cases the identity of the epistemically relevant virtues that agents possess may be called into play, introducing an overly rigorous picture whereby two or more agents are perceived as identical in their cognitive, epistemic or evaluative capacities.

Another related means of idealization, more common in economic theory, is the common prior assumption. While it has encountered a great deal of scrutiny in recent years⁷, the assumption has played a prominent role in the formulation of a dominant conception of rational disagreement between agents. Typically, the attribution of common priors entails that differences in beliefs between agents are explained by

⁷ See (Morris 1995) for an overview.

differences in the information that these agents have⁸. Since by all accounts agents possessing the same priors can only disagree in cases where the information they possess is different. It is thought that in cases of disagreement, updating the common prior with the information that they each possess leads to their each possessing different posterior beliefs (Morris 1995, p.227). But as Kelly has recently pointed out, “this assumption is tantamount to assuming that there is a prior agreement as to the normative import of any piece of evidence which might be encountered” (Kelly, in: Gendler & Hawthorne 2008, p.176). This is an important point. One that shall serve us in showing the difference between a conception of peerage delimiting a domain of expected agreement and an overly abstract conception of peerage assuming agreement with regard to evidence encountered within that domain. It is to a discussion of this last point that I now turn.

Fragmenting conceptions of evidence across epistemic domains

A sceptical approach regarding the possibility of broad and encompassing consequences and generic normative entailments arising from peer disagreement stems, in its simplest form, from the equivocation of notions central to the formation of the normative challenge of disagreement and its needy solution. Notions such as ‘peer,’ ‘evidence’ and ‘fact’ seem to have varying levels of precision in the domains within which disagreement occurs. For instance, the definiteness of the provinces of

⁸ For a characterization of the assumption of common priors see: (Joseph Y. Halpern 2000). See also Robert Aumann’s seminal “agreeing to disagree” for an account that shows that agents with common priors could not agree to disagree (Aumann 1976). As T. Kelly rightly states, the conception underlying the literature stemming from Aumann’s work on disagreement is that “the discovery that another person holds a view that one is inclined to reject constitutes evidence that the other person has access to relevant evidence which one does not possess” (Gendler & Hawthorne 2008, p.176).

consideration in algebra, which may be rather awkwardly referred to as the ‘evidence’ of algebra, is different to the role that evidence plays in cardiology or politics. An algebraic equation, for instance, *entails* a solution by logical necessity and thus when the equation is considered (rather awkwardly) as ‘evidence,’ disagreement with regard to its solution not only has a clear path for resolution but also serves as a flawless indication that at least one party in the disagreement is wrong⁹. This is what makes disagreement in areas such as mathematics *tighter*, and thus stricter in its regulations for revision. While with regard to a cardiac condition there doesn’t seem to be a singular conception of evidence entailing only one possible cause from the given evidence, as there is in algebra. Nor is there any derivative logical law from which such a cause can be unequivocally inferred. So too is the case in politics, where the scope (and limits) of evidence is much harder to demarcate. And where nothing about any given piece of evidence logically entails one and only one interpretation of it.

Therefore there is reason to think that the means of establishing peerage in a given domain sanction different normative regulations for disagreements between peers within that domain, depending on how ‘tight’ the symmetry credentials for peerage can be. The more precise and measurable the epistemically relevant correlation between agents, the more the disagreement between them poses a problem in need of resolution.

⁹ Elga voices this position when objecting to spinelessness: “... in the case of the clean cases [of disagreement] one is in a position to count one’s associates as peers *based on reasoning that is independent of the disputed issue*. But in the messy real-world cases, one is rarely in a position to do so. That is because in the messy cases, one’s reasoning about the disputed issue is tangled up with one’s reasoning about many other matters.” See: (Elga 2007, p.492) and (Pettit 2006) cited by Elga.

I believe that a formulation such as this provides a more feasible answer to the normative problem raised by disagreement. Since it allows for cases where peerage between agents is tighter as well as for cases where peerage is established on less robust, quantitative or mathematical foundations. Of course this is not to say that in less stringent areas of life disagreement between peers does not call for any sort of response, modification or revision. In many cases sticking to your guns when someone just as qualified as you disagrees' may be a visible mark of unlawful persistence or irrational behaviour¹⁰. After all, whether or not you know for certain that someone is wrong – as in algebra – or whether you think someone might be wrong – as in politics – it seems that accommodating the belief of a disagreeing peer in some way, or at the very least not ignoring it altogether, is the reasonable thing to do.

Peerage as the expectation of agreement within a domain

In light of the above I am lead to suggest a refined construal of epistemic peerage according to which the notion is maintained while the conditions for its fulfilment are softened so as to encompass its application across a wide array of domains. Accordingly, two or more people are epistemic peers in a specific domain when we, and they (assuming they recognize that they are peers), have the same *expectations* of each of them being correct in that domain:

EP: two people are epistemic peers with regard to some domain D, inasmuch as they are each *expected* to respond in the same way to new evidence in that

¹⁰ I am grateful to Robert Aumann for pointing this out.

domain. They are expected to respond “in the same way” in that all else being equal, they are expected to agree on whether proposition p is entailed by E .

I see EP as being both an adequate description of instances where peerage is of normative use, while also setting forth the normative entailments that ought to derive from it. What this softer approach can be taken to show is that ordinarily one’s epistemic peers are expected to respond similarly in the domain in which peerage is established. In the sense that if you and I are peers we expect each other to be equally likely to give the right answer to a question in the domain - or in relation to the subject or topic - within which we are peers. This working definition of peerage seems intuitively correct. David Enoch has also expressed it in a recent paper:

“Suppose you trust someone – call him Adam – to be your epistemic peer with regard to a certain topic, for instance philosophy. If asked to evaluate the probability of you giving a correct answer to an unspecified philosophical question and the probability of Adam doing so, you give roughly the same answer. You treat Adam as your philosophical peer. ” (Enoch, p.1). ... [elsewhere:] “by your ‘peer’ I will understand someone who is, somewhat roughly, antecedently as likely as you are to get things right (on matters of the relevant kind)” (Enoch, p.4).

I concur. So this is what seems to be expected of two people who are peers. And the reason it is expected of them and not for instance guaranteed, is that despite their peerage, there is margin for error in each of their responses within this domain. Let us after all remember that the problem of peer disagreement stems from our epistemic imperfection, located in the wider context of human fallibility. And hence margin for error – however small – is expected from our conduct even in those domains in which

prior probability of our being right is adequately established. This is a point that I shall discuss in the following section.

What it is that establishes the expectation of mutual agreement between peers in a domain is the prior evidence for its existence within that domain. As discussed above, typically such evidence comes in the form of competence and credentials for expected judgements within the domain. These are established by past behaviour or performance within some domain D , or within affiliated and relevant domains such as D_1 , D_2 , and D_3 , on the basis of which future performance in D can be inferred. Additionally, expected judgements are established by the unvarying epistemic access to evidence within D . In this latter sense neither peer is thought to be more competent, experienced or privileged in judging or evaluating any $p \in E$ in D .

Thus once a and b are established as peers within D , we expect them to stand equally (have the same attitude) in relation to any proposition p with regard to evidence E within D . We might phrase this form of expectation in probabilistic terms by stating that a and b are epistemic peers in D and according to their past credentials this entails that they each have some particular probability of being correct on any matter p within D . And because they are peers, their probability of being correct on any matter p within D is the same. So that if D were a hat containing any number (n) of p 's, and we were to pick out some p at random, the probability that a and the probability that b being right about p would be the same. Note that because these are non-idealized cases, the probability of being correct in any domain D will be < 1 .

So as to flesh things out let us say that after analysing a , and b 's competence and credentials, each is attributed a 0.8 probability of being correct on any matter within

D (on any $p \in D$). It is their equal probability of being correct which makes a and b peers within D.

Sure enough, a and b needn't be the only peers in D. There may very well be another pair of peers in D, let us call them c and d , who are epistemically superior to a and b , and it is expected that they will be correct on matters within D not at probability 0.8 but rather at probability 0.9. In such a case, c would stand to a not as peer, but as a *superior*. In that we (and a) would expect c to be more likely (+0.1 more likely) to be correct on any given matter within D.

Now if, as in idealized conceptions of peerage, we knew that a and b were identical in all the epistemically relevant factors such as cognitive ability, past experience, success in prediction, equal perception of evidence, and similar evaluative abilities (that is, identical in the sense of their having common priors), then we would be able to say not only that a and b are expected to provide an equally correct answer with regard to any question within D, but also that with regard to *any specific question* within D, they are peers. Meaning that in such a case they are both expected to come up with the right answer for *any particular* question within D.

The difference I am pointing to here is that between two agents being peers in a particular domain and two agents being peers in relation to a particular question within that same domain. As I hope to show, what I will refer to as *absolute peerage* may apply not only to expected responses within a domain but also to expected responses to each and every question within that domain. While the softer conception of peerage I have been advancing, that accounts for cases of non-idealized epistemic imperfection, will apply when peerage entails symmetrical expected responses within

a domain while it doesn't necessitate that this same form of peerage apply to specific questions within that domain.

Absolute and Soft peerage

For two agents to be epistemic peers they needn't have absolute certainty in being correct on any matter in D. They may have a 0.8 likelihood of being correct about matters in D. And this is true of both what I would like to refer to as 'absolute' and 'soft' peers. The difference being that if these are *absolute* epistemic peers, their level of certainty applies to the same 80% of cases in D. So if domain D has 100 possible questions belonging to it, *absolute peers*, just like *soft peers*, are expected to answer 80 of these correctly. But in the case of *absolute peers* they will answer correctly on the same 80 questions.

Absolute peerage

Note that in a case such as this I am describing symmetrical circumstances in which there is no room for any of the peers to err. Saying that *absolute peers* will answer the same 80 questions correctly is tantamount to seeing the domain of their agreement as a predetermined sub-domain of D in which they have *absolute* symmetrical expectations. Thus in the above example, these 'fixed' 80 questions which they are both expected to 'know', can be thought of as D_1 , a sub-domain of D, where $D_1 \subset D$. And where questions in D_1 are immune to error for such absolute peers. For the remaining sub-domain of D, that which is equal to $(D - D_1)$ and is composed of the 20 remaining questions, the absolute peers are expected to err 100% of the time. Thus according to this conception of absolute peerage, within D both agents are expected to

behave exactly the same. They are peers in that they have equal prior credentials (common priors) entitling them to the same *expected* level of credence in D. Furthermore these prior credentials establish them as peers with regard to *any question* within D. In that prior symmetry and entirely equivalent expectations renders their response to any question within D with the same level of credence. Clearly such absolute peerage is an idealized abstraction and it is hard to believe that such peers can actually exist. Let us therefore explore a softer form of peerage. Let us do so while keeping in mind human fallibility and epistemic imperfection, whereby two agents still each have the same probability of success within a domain while this probability does not impinge on specific questions within the domain.

Soft peerage

Let us now envision two peers with equivalent *soft* credentials - relevant cognitive ability, relevant past experience, equivalent success in prediction, similarly relevant ability to perceive and evaluate evidence - entitling us to expect that they will be equally likely to succeed on any question within D. Let us say that this prior data about them establishes their peerage at a predicted 0.8 success-rate with regard to any question within D. Now what this can be taken to mean is that they are equally likely to be correct on any question in D; in fact, they are 0.8 likely to be correct on any such question. And yet all the same, with regard to any specific question p within D, we do not know and cannot establish whether p falls within their common area of peerage.

As a point of case let us say that prior to a random question draw from within D, persons' a and b , who are (*soft*) epistemic peers within D, share an equal likelihood of being correct on any single question randomly drawn out of a hat containing all the

questions in D. Yet once the question has been drawn we cannot establish whether *a* and *b* are epistemic peers with regard to *that* specific question. Since nothing about *a* and *b*'s prior credentials can establish this form of expectation. There doesn't seem to be any prior information about *a* or *b* that can tell us anything about how they both stand toward the specific question that has been drawn. All we know is that they are equally likely to be correct within the domain as a whole.

Being peers on subjects and questions: the case of epistemic imperfection

Let us elaborate for a moment on the idea of peerage with regard to a question within a domain as opposed to peerage with regard to a domain. Two agents are epistemic peers with regard to a given question if prior information about them grants they possess equal probability in answering the specific question correctly. Thus if our domain is *Cold War history*, a question within that domain may relate to 'Eisenhower's attitude toward Nasser's plans concerning the Suez Canal.' In this case epistemic peerage with regard to *this* question must be established by credentials ensuring that *a* and *b* respond in the same way (or that there is no reason to think that either of them has any privilege or superiority with regard to answering this question). But as a matter of fact, we do not know if *a* and *b* are in fact equally likely to answer this question correctly. This question may after all reside in the (D - D₁) subdomain of their peerage in D, a sub-domain of D in which they are not peers. Consider that we know *a* and *b* are 0.8 likely to be right in D. Hence the likelihood of their beliefs overlapping and being right in D is 0.64. But this also leaves an area where there beliefs do not overlap and hence it leaves an area in which they are not in fact peers.

Another example would be if we say of Pat and Harry that they are epistemic peers in the prediction of inflation (Pat and Harry are economists). They have an equal success

rate of having done so in the past, they have the same credentials and qualifications - they have worked in equally ranked firms at parallel positions, with an equal past track record of achievement and evaluation. While we can say of Pat and Harry that they each have an equal probability of being correct in their prediction, we cannot be as sure that they each have an equal probability of being correct with regard to a specific prediction – say the inflation rate in Egypt following the January 2011 riots. It doesn't seem that anything about Pat or Harry can tell us that they are equally likely to be correct on *this* question. As far as we know, Pat may be more versed in middle-eastern affairs or perhaps just more likely than Harry in being correct with regard to interpreting the evidence of this specific state of affairs. But this is not a point about what Pat *is* more competent about but about what she *may* be more competent about.

Above, so as to differentiate the idealized cases of epistemic peerage from disagreement in its more familiar surroundings, I gave an example in which the past history of two agents established a 0.8 credence ranking in a domain in which they are both peers. But now we can see that so long as this level of credence is below full certainty, i.e., below probability 1, then even if we know that two agents each have a 0.99 chance of being correct on any question p within D , so long as these are *soft* peers, we can never know whether this question falls in the 99 percentile and thus whether it is a question in relation to which the two agents are peers or one which acts as a peerage-breaking-question, where the disagreement itself provides one with a reason for the demotion of one's peer.

Thus if Pat and Harry are epistemic peers each with a 0.99 probability of being correct on any question within D , assuming the number of questions in $D \geq 100$, we still cannot say that they are epistemic peers on any given question within D . Since we cannot be sure that any question chosen is not a question that exceeds the scope of

their peerage. And thus upon disagreement with regard to some question p , both Pat and Harry can claim broken-peerage as the reason because of which they now disagree. Similarly, as neutral third-party observers, we too cannot be fully justified in claiming that Pat and Harry are peers on any given question, but rather only in the domain in which the question resides.

The *soft* account that I have presented maintains the commonplace normative significance of the notion of peerage while also pointing to its fragility in cases of disagreement. Therefore it may be rational for us to have different bets prior to the draw of a question from D and after the question has been drawn¹¹. Thus we should pay attention to the rationale for placing bets *ex ante* and for placing bets *ex post*. For the same reason, the probability I grant my peer being right on a certain matter need not be equal to my probability of his being right given a disagreement between us with regard to that matter (Enoch, p.24). And assuming that I recognize this difference, my prior probability (prior to the draw of a question from the D -hat) should reflect my somewhat cautious expectation of our mutual agreement when the time comes.

Experts and subdomains

I began by arguing for a differentiation between the normative use of peerage and the means by which this notion is established. I showed that what establishes peerage is evidence in its support. I showed that the nature and rigor of evidence for peerage varies across domains. I also claimed that there is a correlation between the nature of the evidence establishing peerage and the normative implications which stem from

¹¹ Just like in the Monty Hall Problem where the player who swaps his initial choice of door has a $2/3$ chance of winning while the player who sticks with his initial choice has a $1/3$ chance of winning.

acknowledged disagreement between peers. I differentiated between an absolute conception of peerage and a soft conception. Whereby according to the soft conception peerage was established on grounds of relevant epistemic credentials, ensuring that while peers may have the same likelihood of being correct on questions within a domain, their peerage, or shared expected likelihood, does not extend to each and every question within that domain. As I showed, it is on this point that *soft* and *absolute* peers differ. Soft peerage, which aims to more adequately characterize real-world forms of peerage and their inherent fallibility, leaves open the possibility of peerage-breaking questions within a shared domain of equal expertise. And this also leads to an additional possibility that isn't enabled by the absolute conception of peerage.

Two weather forecasters, *a* and *b*, justifiably¹², see one another as peers when it comes to weather forecasting. But they also acknowledge that there may be areas within their domain of shared expertise in which one is better than the other; areas within weather forecasting in which *a* is better than *b* and areas where *b* is better than *a*. In this sense we expect *a* and *b* to be as good as each other when it comes to weather forecasting in general, and they too expect this of each other. But all the same neither we, nor they, expect them to be identical in all their predictions.

Thus within the domain in which we are peers, there *may* for instance be questions that are of a type at which I am better and questions of a type at which you are better.

¹² See (Enoch, sec.5) for a discussion of the normative significance of 'being a peer' vs. 'being believed to be a peer,' and 'being justifiably believed to be a peer.' Enoch's claim is that the question of how to revise beliefs in light of a discovered disagreement is actually ambiguous between three readings (Enoch, pp.20-25).

We may be aware of these types of questions, but this is not necessary. This is something we might assume as possible and validate post factum. But notwithstanding the possibility of our each being better on questions of a certain type, in general there are an equal number of both types of questions in our shared domain of expertise. Hence with regard to any question that would be randomly pulled out of a hat containing all the questions in our domain of peerage, neither of us would have any reason to think himself more likely to be right or any less likely to be wrong¹³. Hence this - rather reasonable scenario - allows that the probability that we each have for being correct on a given question within the domain is the same, while we know that on any given question chosen at random we do not share a common likelihood for being right.

Bibliography

- Aumann, R.J., 1976. Agreeing to Disagree. *The Annals of Statistics*, 4(6), 1236-1239.
- Christensen, D., 2007. Epistemology of disagreement: The good news. *Philosophical Review*, 116(2), 187-217.

¹³ Note that the possibility of our domain of shared expertise having sub-domain in which we are each experts is not present in the formulation of the problem of disagreement in the literature. It is an addition I have put in place as a means of relaxing the constraints on peerage. According to this additional point, the fact that as peers we do not know whether the new question is within our subdomain of expertise maintains the tension of what we ought to believe now (given that we should recede to our prior conditional probability) but all the same it allows for it to be possible that as peers it is not necessary that we be identical with regard to everything in our shared domain.

- Christensen, D., 2010. Higher-order evidence, *Philosophy and Phenomenological Research*, 81(1), 185-215.
- Elga, A., 2009. How to disagree about how to disagree. In T. Warfield & R. Feldman, eds. *Disagreement*, Oxford University Press.
- Elga, A., 2007. Reflection and disagreement. *Noûs*, 41(3), 478–502.
- Enoch, D., Not Just a Truthometer: Taking Oneself Seriously (but not Too Seriously) in Cases of Peer Disagreement. *Mind*. Available at: <http://mind.oxfordjournals.org/content/early/2010/12/21/mind.fzq070.abstract> [Accessed February 9, 2011]
- Feldman, R., 2006. Epistemological puzzles about disagreement, In S. Hetherington, ed. *Epistemology Futures*, Oxford University Press, pp. 216-326.
- Feldman, R., 2007. Reasonable religious disagreements, In L. Antony, ed. *Philosophers Without Gods: Meditations on Atheism and the Secular*, Oup, pp. 194-214.
- Feldman, R. & Warfield, T.A., 2010. *Disagreement*, Oxford University Press.
- Gendler, T.S. & Hawthorne, J., 2008. *Oxford Studies in Epistemology*, Oxford University Press.
- Joseph Y. Halpern, 2000. *Characterizing the Common Prior Assumption*, EconWPA. Available at: <http://ideas.repec.org/p/wpa/wuwpga/0004009.html> [Accessed February 7, 2011].
- Kelly, T., 2008. Disagreement, Dogmatism, and Belief Polarization, *The Journal of Philosophy*, 105(10), 611-633.
- Kelly, T., 2009. Peer disagreement and higher order evidence, In T. Warfield & R. Feldman, eds. *Disagreement*, Oxford University Press.

- Kelly, T., 2005. The epistemic significance of disagreement. In J. Hawthorne & T. G. Szabo, eds. *Oxford Studies in Epistemology, Volume 1*. Oxford University Press, pp.167-196.
- Morris, S., 1995. The Common Prior Assumption in Economic Theory. *Economics and Philosophy*, 11(02), 227-253.
- Pettit, P., 2006. When to defer to majority testimony - and when not. *Analysis*, 66(291), 179-187.
- Rosen, G., 2001. Nominalism, Naturalism, Epistemic Relativism. *Nous*, 35(s15), 69-91.