

Overcoming Expert Disagreement In A Delphi Process. An Exercise In Reverse Epistemology

Elisabetta Lalumera[†]
elisabetta.lalumera@unimib.it

ABSTRACT

Disagreement among experts is a central topic in social epistemology. What should an expert do when confronted with the different opinion of an epistemic peer? Possible answers include the steadfast view (holding to one's belief), the abstemious view (suspending one's judgment), and moderate conciliatory views, which specify criteria for belief change when a peer's different opinion is encountered. The practice of Delphi techniques in healthcare, medicine, and social sciences provides a real-life case study of expert disagreement, where disagreement is gradually transformed into consensus. An analysis of Delphi shows that moderate conciliatory views are descriptively more adequate than rival views. However, it also casts doubt on whether the debate in social epistemology is explanatory relevant vis-à-vis real life cases of expert disagreement, where consensus replaces truth, and acceptance is more explanatorily relevant than belief.

Keywords: disagreement, experts, epistemology, Delphi process, acceptance.

1. Introduction

A central topic in social epistemology is the problem of disagreement among experts. What is rational for an expert to do when one of her peers does not share her opinion? Answers to this question that have been provided include the “abstemious view”, according to which suspension of judgement is the only rational option when one acknowledges her peers' dissent (Feldman 2006, 2007; Christensen 2007), the “steadfast view” (one should hold fast to one's beliefs, dissenting epistemic peers notwithstanding) (Goldman 2010b, van Inwagen 1996, Kelly 2005), and moderate conciliatory positions aimed at

[†] University of Milano Bicocca, Italy.

specifying conditions in which belief revision is appropriate (Elga 2010, Lackey 2010, Christensen and Lackey 2013).

The problem of disagreement about experts is also a central topic in the life of scientific communities, governments, and organizations in general, where panel of experts are called in to report or even to decide on technical issues. Here the problem takes the form of reaching a final verdict in spite of the possible and actual disagreement within the panel of experts. The present paper discusses a widespread method for reaching consensus among experts, the Delphi technique. Introduced in the 60s, Delphi is now routinely employed in many areas, typically medical and healthcare services research, but also social sciences, in order to receive answers from expert peers who are likely to disagree, but whose consensus on a certain topic is important to reach (Dalkey and Helmer 1963; Linstone and Turoff 1975). This paper is an exercise of reverse epistemology in that the question “what should experts do when they disagree?” is confronted with “what do experts do when they disagree?” and the guiding principles of what is considered a good practice of overcoming disagreement are abstracted out and assessed, in a bottom-up rather than top-down approach. The guiding idea here is that the logical analyses that classical normative epistemology provides, can be tested with, or employed for real-life cases of knowledge formation procedures, and real-life problems of disagreement. This is not a widespread attitude in the literature yet, though Boaz Miller tested real-life cases with classical epistemological solutions to the novice-to-expert disagreement problem (Miller 2014), and Nathan Ballantyne suggests that the new challenges for classical epistemology crucially involve moving away from oversimplified, toy examples (Ballantyne 2014, see De Cruz and De Smedt 2013 for dissent).

The paper advances the following conclusions. On the assumption that the Delphi procedure is an *epistemically rational* practice, and that it can be considered a practice of *knowledge* formation, it appears that moderate conciliatory views are descriptively more adequate than the steadfast view and the abstemious view. However, Delphi techniques qualify as rational epistemic procedures vis-à-vis expert disagreement only if we concede that consensus - and not just knowledge - is an epistemic value in itself, or alternatively that consensus can be knowledge-based. Also, Delphi consensus practices involve the pragmatists' claim (from William James) that acceptance - and not just belief - is a relevant mental state in assessing knowledge and knowledge procedures.

2. Expert disagreement in classical epistemology

In this section I provide an overview of the main positions on expert disagreement in classical epistemology – the philosophical study of what should we do to achieve knowledge, conducted with the method of conceptual analysis. Differences among authors aside, classical epistemology is a normative, rather than descriptive discipline. Knowledge, the main epistemic goal, is generally conceived as connected with truth and justification – rather than as the mere acquisition and systematization of information (Goldman 2010a).

Within this tradition, the problem of expert disagreement is a particular case of the more general problem of disagreement among *epistemic peers*. Peers can be characterized as subjects who are as well-informed and well-disposed to react to the evidence as we are. This characterization stresses parity of intellectual and cognitive virtues, as well as shared evidence (Gutting 1982, Kelly 2005, Christensen 2007, Lackey 2008). A slightly different way to define epistemic peerhood involves expected outputs: a peer is someone who we think is just as likely to be mistaken in her judgement, as we ourselves are (Elga 2007). On this latter view, one's antecedent epistemic score, so to say, is relevant to peerhood. In the typical toy example often discussed in the literature – the simple maths case - my friend and me are plausibly epistemic peers when we confront with the problem of checking the bill that the waiter brings us at the restaurant after dinner, provided that we are equipped with same mathematical skills, same evidence, and in absence of defeating conditions of perception, cognition, and other physical abilities that may be relevant (Christensen 2007, Elga 2007).

Experts are a special kind of epistemic peers. Following Alvin Goldman's definition, an expert is "someone who possesses an extensive fund of knowledge (true belief) and a set of skills or methods for apt and successful deployment of this knowledge to new questions in the domain. Anyone purporting to be a (cognitive) expert in a given domain will claim to have such a fund and set of methods, and will claim to have true answer(s) to the questions under dispute because he has applied his fund and his methods to this question" (Goldman 2001, 146). Experts are not epistemic peers with respect to novices and non-experts, but they can be epistemic peers of other experts. In fact, expert communities and panels are (ideally) groups of people who are both experts with respect to novices and the general public, and peers among

themselves, that is, each of them ideally thinks of the others as being just as likely to be right as he or she is, in judgments pertaining a given area.

We can imagine sets of epistemic peers located at any level of competence and epistemic status, peerhood being a horizontal relation of equivalence holding between subjects (independently of the quality of the evidence or competence they possess). Novices can be peers, and experts can be, with respect to a particular judgement (Elgin 2010). Thus, the problem of expert disagreement is just a special case of the general problem of disagreement among peers, which can be stated as follows: Once that we realize that a peer does not share our belief (or belief confidence) in the face of the same evidence, what should we do? Do we have at least one reason (the peer's disagreement) to revise it, or rather we should withdraw judgement, or hold fast to our view? In the simple maths example, what should I do when I realize that adding up the bill I arrive at 46 euros, while your result is 48?

The debate on peer disagreement is now characterized by a continuum of positions ranging from Strong Conciliationism to Steadfast positions (Christensen 2009, Christensen and Lackey 2013). Conciliationist views in general hold that a peer's disagreement always constitutes for a rational subject a good reason to undergo a process of belief revision. When a disagreement about a certain proposition p is revealed, everyone should give some weight to her peer's judgment such that neither is justified in staying exactly as confident as she was before regarding p . The intuitive motivation behind Conciliationism is that to recognize someone else as an epistemic peer just is to assume that we both have the same chance of being wrong, and nothing is special (epistemically speaking) about me. According to Conciliationist, assuming peerhood makes one irrational if she doesn't modify her attitude upon disclosure of disagreement (Feldman 2006, 2007, Christensen 2007, Elga 2010, Ballantyne and Coffman 2012, Carter 2014).

One way – the strongest way - to reduce confidence in one's belief when disagreement is faced, is to suspend one's judgment altogether. This is what Richard Feldman famously argued for:

One of us must be making some kind of mistake or failing to see some truth. But I have no basis for thinking that the one making the mistake is him rather than me. And the same is true of him. And in that case, the right thing for both of us to do is suspend judgment on p (Feldman 2007, 212).

This particular variety of Conciliatory view is now called the Abstemious view, for the idea is to abstain from judging about p , when your epistemic peers disagree (Aikin et al. 2010). As Feldman and Kornblith noted, it seems that suspending judgement is intuitively appealing in some cases, but not others. Perceptual cases are of that sort - if peers A and B look out of the same window from the same angle and A sees a pidgeon on a tree branch, and B does not, it seems rational for them to suspend judgement over ‘there is a pidgeon on that branch’. Abstention seems appealing in this kind of case because, assuming peerhood, the epistemic viewpoint of the disagreeing peers looks like a dead end, there is no room for checking again, and try to locate the possible mistake. A and B may try again to establish whether there is a pidgeon on the branch, but they would better do that through a different epistemic route, changing the evidential basis or the procedure. For example, A may find out a telescope, and B could choose another angle from which to look at. The same goes with the simple maths case I mentioned above. If two peers at the restaurant disagree over the bill, it is rational for each of them to suspend their judgement, and change the epistemic settings completely, by using the mobile calculator, or asking a third party to check (Feldman 2006, Christensen 2007, Kornblith 2010).

There are two characteristics of the Abstemious position, which are of special relevance with respect to the concern of this paper, namely, an application of epistemological questions to a real-life epistemic practice such as Delphi. The first one is its relation with the Uniqueness Principle (Feldman 2006). As some philosophers have pointed out, Abstemious views on disagreement are appropriate if a Uniqueness Principle about evidence applies, stating that for any given proposition p and body of evidence, the evidence fully justifies just *one* level of confidence in the proposition (White 2005, Kelly 2005). Simply stated, the idea is that two peers A and B that disagree over p , can’t be both right. Evidence cannot be interpreted in more than one way. Therefore, if neither A nor B has reason to claim a better epistemic status than the other one, suspending judgement over p comes out as a reasonable option. In the pidgeon-in-the-field perceptual case, and the simple maths case, Uniqueness seems appropriate, as the evidence is per hypothesis complete and sufficient for establishing whether p is true or not: adding numbers just gives one result, and checking perception with reality just gives one verdict. In Earl Conee’s terms, “we have no better basis for discounting opposing summary impressions than we do for our own” (Conee 2009, 315).

The second characteristic of the Abstemious positions is that it leads to widespread skepticism, if for most (or maybe for all) possible opinions that p that one has, one can find an epistemic peer who disagrees, then for most (or maybe for all) possible opinions that p , one cannot say one knows whether p or not p . As Christensen puts it, there must be something wrong with a view that takes massive suspension of belief as the right thing to do (Christensen 2009). As other kinds of skepticism, Conciliatory Abstemious views on peer disagreement put too high a threshold to what counts as rational belief: in areas where disagreement among peers is pervasive (including politics, morality, arts, and science), almost nothing of interest would count as a rational belief. In fact, the typical cases (the simple maths and the perceptual case) that fuel intuitions in favour of this view are such that the value of having a belief at all, and reaching a verdict in that particular situation, is close to none (the bill can be easily checked later, or with a calculator, and nothing depends on whether there really is a pidgeon on the branch). Kornblith (2010) adds that to abstain from believing is more plausible, and more rationally defensible, when the beliefs in question are disconnected, or not much connected with the others one holds. If suspending my belief that p implies suspending my endorsement of a whole theory that is logically entailed or presupposed by p , then abstention is in conflict with an intuitive principle of conservation. Beliefs that God does not exist, or that the Earth is not flat – unlike “there is a pidgeon on the branch” are examples of beliefs whose withdrawal or suspension in the face of disagreement would cause a huge revision of a person’s overall system of convictions.

The above problems of abstemious views, and the considerations of different kinds of examples, speak in favour of more moderate forms of Conciliationism. The common idea is that acknowledging one’s peer disagreement makes it rational to change one’s belief or level of confidence, not to withdraw one’s judgment. One simple way to go is to reduce one’s level of confidence in one’s judgment by “splitting the difference”. Thus, for example, experts A and B assess evidence, A reaches the judgment that p with credence 0.8, and B reaches the judgment that p , but with credence 0.2; then they learn about each other’s verdicts, and compromise by both judging that p with credence 0.5, the average of one’s own and the other’s credence. Conciliationism of this sort is epistemically fair and founded on the principle that one’s own judgement and reasoning and one’s peers’ judgment and reasoning have equal weight (Elga 2007). The possibility of each of the peers’

being completely irrational or wrong in that particular occasion makes this strategy epistemically risky – if B were just badly wrong in assessing the evidence, A's lowering her credence to 0.5 would be incorrect (Kelly 2005). So now Conciliationists tend not to adopt "splitting the difference" as a rule of thumb: a peer's disagreement is just one piece of evidence to be considered among others (Christensen 2009). A possible correction comes from complicating the examples: if disagreeing experts are more than one, then A's compromise is less risky, and more rational. Jennifer Lackey (2008) has proposed another sort of correction: a peer should compromise to other peer's disagreeing judgment depending on her level of "justified confidence" – the more A is justifiably sure about believing that p , the less she should compromise. Suppose (Lackey's own example) that A and B are doctors, and p is "patient C suffers from lupus", where patient C shows many characteristic symptoms but no skin rash. Same evidence and level of expertise could still make room for doctor A being certain only to 0.5 credence, because she has never seen a patient with lupus with no skin rash. In that case it would be rational for doctor A to compromise to her peer's opinion.

At the opposite end of the spectrum with respect to Abstemious views, and in reaction to the difficulties of Conciliationism, classical epistemological positions about peer (expert) disagreement feature Steadfast views. According to the Steadfast view, when one expert discovers that another expert (or other experts) disagree over her verdict about p , she is still rationally permitted to hold fast to her belief (Goldman 2010b, van Inwagen 1996, Kelly 2005). The common strategy employed in favour of this view is to deny that one's personal point of view is the same as the other peers' point of view – in sum, to deny some aspects of peerhood, while granting sameness of evidence, competence and expertise. One straightforward way is to admit that vis-à-vis the same evidence (and granting the same competence and expertise), there can be different epistemic norms, i.e. methods, procedures, and ways to measure the importance of evidence that different agents can employ. Epistemic norms say that some sources or ways of forming a belief are reliable, i.e., likely to produce true beliefs, and examples may vary from particular diagnostic procedures in medicine, to meta analyses, genetic research, use of animals in testing, computer simulations, integration of different kinds multiple lines of evidence.

This is a form of epistemic relativity: what counts as rational, and rationally known, may vary as standards vary. Thus, expert A arrives at judging that p operating with norms N , then she might well realize that expert B arrives at

judging that not p operating with norms M , but given that A does not endorse norms N , she is rationally permitted to hold fast to p (Goldman 2010b). Epistemic relativity is incompatible with the Uniqueness Principle I mentioned above – that certain evidence commands only one possible verdict.

There is also a second way by which one may come to deny peerhood, and support the Steadfast view on disagreement, which is slightly similar to the first one, but more focused on a conceptual feature of beliefs and acts of judgment. It is the idea that even though one may admit or know that another person's reasoning is sound, and her principles are good enough, one's act of judging can be moved by one's own reasons only. So, for example, Bergmann (2009, 339) and Elgin (2010) argue that expert philosopher A may still hold her conviction that p even though she fully understand the reasons of expert philosopher B , who denies p , because understanding reasons is not the same as endorsing them. One's own reasons guide one in a way that other people's reasons do not, more precisely, one's epistemic reasons lead one irresistibly to form a belief given a certain evidence (Wedgwood 2010). Elgin on this point cites Bernard Williams' point that belief is something that happens to a subject independently of, and in spite of, her will or intention – it is not voluntary (Williams 1973). Thus, she argues, "since beliefs are not voluntary, an epistemic agent cannot, even through judicious assessment, bring it about that she retains, lowers her degree of belief, or suspends belief in the face of a disagreement (Elgin 2010, 12)". Thus, disclosure of disagreement or even a report of my peer's own reasons for disagreeing are sufficient to move a belief change in me; only if the other person's reasons become my own reasons (if I enlarge my evidential basis so to comprise them), they can come to affect what I believe. Note that this point is against any form of Conciliationism: by acknowledging another person's reasons no one can *ever* change her beliefs.

There is, however, an interesting turn in Elgin's argument against Conciliationism. She claims that when an epistemic agent acknowledges her peers' disagreement, and their reasons, she may, however, be able to affect her responses" (Elgin 2010, 12). How? By *accepting* other peers' judgment, where acceptance is a different mental state than belief (while possibly retaining the original belief). The difference between belief and acceptance is usually put in these terms: while belief is involuntary, acceptances are voluntary actions, so that one can accept that p at will, for practical or prudential reasons. To accept that p is to commit oneself to adopt it as a premise in inferences, or as a basis for action, as if it were true (Cohen 1992,

Wray 2010). One can accept that a certain person is trustworthy and capable – even without conclusive evidence – because she needs to employ her with a specific role in a plan of action. One can accept a scientific hypothesis and treat it as a premise, to see what follows, and check it. Acceptances are the attitude we bear to the antecedents of conditionals, when we reason hypothetically. The interesting point here is that if Elgin and Wedgwood are right, peer disagreement should affect what we accept. It should influence the inferences we are prepared to make and actions we are prepared to perform. I will say more on this in the next section.

3. How experts overcome disagreement in a Delphi process

This section confronts the positions in classical epistemology about expert disagreement, summarized above, with the Delphi process, a methodology aimed at issuing guidelines, advices, forecasts and consensus statements in general, in cases and fields where experts are likely to disagree, and/or purely statistical methods of analysis of the data are not possible, or impractical. Introduced around 1965 as a procedure intended to “obtain the most reliable consensus of opinion of a group of experts . . . by a series of intensive questionnaires interspersed with controlled opinion feedback” (Dalkey & Helmer, 1963, 458), the Delphi process is now widely employed in a variety of domains, prominently including medicine, healthcare, and the social sciences, but also finance, engineering, geography, geology, and many others (Gupta and Clark 1996 for a review, Powell 2003 for a review of criticisms and shortcomings of the method).

A Delphi process involves (at minimum) the following steps: definition of a problem for which a judgment, or a guideline, is required; selection of a panel of experts; definition of a series of questions that specify aspects of the problem, usually but not always formulated or assessed by the experts themselves; a first round of questionnaire, where each expert is sent or confronted with the questions so that she can give her answer anonymously; statistical analysis of the results; publication of results of the first round questionnaire; optional iteration of questionnaire; and final verdict. Here is an example of a Delphi process about the diffusion of Alzheimer’s disease. More precisely, the study was aimed at establishing the value of dementia prevalence for all regions of the world, in 5-year age bands to 84 years, and for those aged

85 years and older (Ferri et al. 2006). Before the study, there was no agreement on such value, partially because of the lack of reliable epidemiological research from some world regions, so a qualitative consultation involving experts was required. As a first step, 12 experts were selected, who had published studies on the prevalence of dementia in peer-reviewed journals. Then, experts were sent the questions about prevalence in each area, together with a document combining the research evidence. Each one sent back his or her own estimates for dementia prevalence in the 14 WHO regions. After this first round, agreement on the value of prevalence of dementia for the region denominated “Afro E” was 0.62; the questionnaire was then sent again, together with the results, distribution of opinions, and other information, so that after the second round the agreement reached 0.80. As the authors write, “members of the panel who were invited to reconsider in the light of their colleagues’ prevalence estimates and any accompanying comments. If they chose to alter their estimates, they could see the effect of this change on the group mean” (Ferri et al. p.3). Finally, the experts’ judgments largely converged, and they were further statistically aggregated so that the final statement was issued: “We estimate that 24 million people have dementia today and that this amount will double every 20 years to 42_{SEP} million by 2020 and 81 million by 2040, assuming no changes in mortality, and no effective prevention strategies or curative treatments. Of those with dementia, 60% live in developing countries, with this number rising to 71% by 2040” (p.4).

A number of observations are in place between tackling the issue of expert disagreement in a Delphi process. First, as a very general point, there is no a priori warrant that the outcome of such processes is knowledge, in the classical sense connected with true and justified belief. Simply put, scientific knowledge is fallible and approximates truth by trying to reach consensus (Miller 2014, Solomon 2007, Steup 2010). Moreover, even excluding culpable intentions, each step, from the selection of experts, to the formulation of questions, to the selection of input evidence given to panelists is prone to error, or may be altered by unconscious biases. The whole tradition of the sociology and philosophy of science following Thomas Kuhn (1970) reminds us that scientists’ personal values, psychological dispositions, and social factors have an impact on their epistemic judgments and on the way they come to issue them. With such disclaimers in place, Delphi processes qualify as rational standards for knowledge acquisition in the enlarged scientific community, and

in this qualified sense they can be taken as a model of how experts overcome disagreement.

As the example shows, at the end the disagreement is in fact overcome. In Delphi processes, sometimes many rounds are needed, but a final judgment is always produced, though sometimes it may be complicated one, incorporating different views – as Miriam Solomon reminds us) (Solomon 2007). Assuming that this method of overcoming peer expert's disagreement is a rational practice, which one of the classical epistemological models describes it more adequately? What happens if we test epistemological models in a bottom-up way?

Let us focus on the transition between low agreement after the first round, and better agreement after the second. Here, Abstemious views and Steadfast views do not seem to capture what is happening. In fact, if all the experts were Steadfast defenders of their prior beliefs, no agreement could ever be reached. With respect to the Steadfast view, I think the problem can be located within the asymmetry first-person claim that the position assumes. A Delphi process is such that at each round experts are confronted with their peers' motivations and evidence for the judgments they gave. As consensus gets formed, each one incorporates at least some of the others' reasons in her own evidential basis or judgment procedure. Ideally, each expert learns in each round. So even though a plurality of possible judgments is assumed to be possible, it is also assumed that the disclosure of someone else's preferred methodology or evidence weighting may alter one's previous positions. Other people's reasons in this contexts are not motivationally inert with respect to one's judgment. The nature of the reasons in questions – scientific reasons – makes the transitions from objective to subjective reasons more feasible. Replicability and objectivity are the key concepts here: scientists are trained to assume and to require that the reasons they employ can and should become other scientist's reasons

Abstemious views are equally inadequate, for the epistemic principle they incorporate is at odds with the characteristics of scientific knowledge formation and theorization. As discussed in the previous section, one of the ways to defend judgment suspension in the face of peer disagreement is the Uniqueness Principle, stating that from a certain body of evidence, only one verdict is mandated. There is no room for underdetermination. Such a principle cannot be assumed without reservation in the domain of science – in scientific research, some underdetermination is generally accepted, implying that different theories and judgments can be compatible with the same evidence, at least at earlier stages of theorization in a given domain. Experts

know that other peers can reach different verdicts, typically if they give different values to certain methodologies and techniques, with respect to others. As said above, the mutual disclosure of specific methodologies and preferred techniques can eventually bring to consensus during the process, but the possibility of initial disagreement of correct verdict is assumed as normal.

So far, I discarded the Abstentionist view and the Steadfast view of expert peer disagreement as adequate rationalizations of the Delphi technique, because they incorporated objectionable principles. There is also another reason why views that recommend judgment suspension are not adequate vis-à-vis the Delphi example. Abstemious views are well-grounded in cases where nothing depends on our suspension of belief. Experts in a Delphi process, on the contrary, conform to William James' claim that "sometimes we are *obliged* to form beliefs on insufficient evidence, and that it would be a significant intellectual, prudential or even moral failure to do otherwise" (James 1896/1979, 298). The nature of this ought can be debated – it can be that consensus is evaluated as desirable and good, as a mark of a mature science. Or differently, the content of that specific agreement could be given a moral value, for example, an expert may believe that she ought to contribute to a consensus judgement about the prevalence of Alzheimer because without that, people who suffer from that illness would not receive appropriate care. Finally – as the sociology of science reminds us - the value of agreement can be of a egoistic, non-virtuous sort: to please a pharma sponsor, or to meet the expectations of the organizers of the survey. Whatever it is, the value of agreement makes it practically rational for each expert to contribute to it, in spite of the initial disagreement.

The practical value of agreement and consensus makes an argument against Steadfast views, as well. If consensus is what experts ought to reach, then a stubborn defence of one's own opinion becomes not valuable in itself, but only provided one's own epistemic reasons are very strong.

Then what is the best way to describe the actual process of belief revision that happens in Delphi rounds? What does each expert do? How is it that agreement is produced? A proper answer of this question would feature empirical psychological investigation. However, a projection from the example described above suggests that moderate Conciliatory positions may well be equipped to provide a correct picture. Given that the distributions of answers and values of credence are disclosed to each member of the panel, it is plausible that processes such as those Lackey envisages are in play: each one

adjusts her own level of credence depending on how high it was antecedently, and on how many peers disagree, or alternatively, peers “split the difference”. According to this description, experts actually change their beliefs, or the level of credence they attach to them, in the face of disagreement, by partially compromising to the other peers’ belief when it is epistemically possible for them to do so.

However, in light of the above considerations about the value of agreement, there is room for an alternative description, according to which acceptance, rather than belief, is the right concept that describes what happens when consensus is reached. Given that the very production of a verdict has a value for all the participants, each one may be described as willing to adopt certain stances in order to converge with other people’s opinions – to hold them as if they were true, in the future practice, even though they still deem the evidence insufficient. For example, if expert A and expert B “split the difference” of their respective credences, and reach a verdict that averages such values, the verdict is voluntarily issued, rather than spontaneously inferred from the evidence. This is to say that in Delphi processes, experts’ rational behavior when facing disagreement conforms to Moderate Conciliationism, but what is produced is actually a state of acceptance.

Note that philosophers have also argued that acceptances, rather than beliefs, are more apt to describe what happens in group decisions, at the collective level (Wray 2010, Miller 2014) – here, I suggest that the Delphi example shows that acceptance may be an appropriate concept to employ also at the individual level, to describe what each expert does when facing disagreement, while holding some interest, drive or desire to overcome it. This is to say that practical rationality and epistemic rationality overlap in real-life cases of disagreement, in a way that is absent from the theorization of classical epistemology.

4. Conclusions

Delphi processes, in which disagreeing experts are brought to agree on a verdict, show that normative accounts of expert disagreement in classical epistemology are mostly descriptively inadequate. Neither Abstemious views, nor simple Steadfast views turn out to be tenable. Of course one may adopt a massive error theory, and claim that scientific practices of consensus are de facto irrational, thereby saving the classical epistemological views as normative

standards. However, such a stance would be implausibly revisionary. Moderate conciliatory accounts of expert disagreement, holding that a disagreement elicits some contextually modulated form of belief revision from the part of each expert, turn out to be illuminating, especially if they are reframed in terms of acceptance rather than belief.

ACKNOWLEDGEMENTS

I would like to thank Maria Cristina Amoretti, Edoardo Datteri, Stefano Fanti, and Paolo Labinaz for their comments on earlier version of this paper.

REFERENCES

- Aikin, S. F., Harbour, M., Neufeld, J., & Talisse, R. B. (2011). On Epistemic Abstemiousness: A Reply to Bundy. *Logos & Episteme*, 2(3), 425-428.
- Ballantyne & Coffman (2012). Conciliationism and Uniqueness. *Australasian Journal of Philosophy*, 90(4), 657-670.
- Ballantyne, N. (2014). Review of *The Epistemology of Disagreement: New Essays*, eds. D. Christensen & J. Lackey, Oxford University Press, 2010. *Notre Dame Philosophical Reviews*, 5 July (online at: <https://ndpr.nd.edu/news/49177-the-epistemology-of-disagreement-new-essays/>, accessed April 14, 2015)
- Carter, J. (2014). Group Peer Disagreement. *Ratio*, online first (online at: [http://onlinelibrary.wiley.com/journal/10.1111/\(ISSN\)1467-9329/earlyview](http://onlinelibrary.wiley.com/journal/10.1111/(ISSN)1467-9329/earlyview), accessed April 14, 2015)
- Christensen, D. (2007). Epistemology of disagreement: The good news. *Philosophical Review*, 116, 187-217.
- Christensen, D. (2009). Disagreement as evidence: The epistemology of controversy. *Philosophy Compass*, 4(5), 756-767.
- Christensen, D., & Lackey, J. (Eds.) (2013). *The epistemology of disagreement: New essays*. Oxford: Oxford University Press.
- Conce, E. (2009). Peccage. *Episteme* 6, 313-323.

- Dalkey, N., Helmer, O. (1963). An experimental application of the Delphi method to the use of experts. *Management Science* 9(3), 458–467.
- Delbecq, A. L., Van de Ven, A. H., & Gustafson, D. H. (1975). *Group techniques for program planning: A guide to nominal group and Delphi processes*. Glenview, IL: Scott Foresman.
- Elga, A. (2007). Reflection and Disagreement. *Noûs* 41, 478–502.
- Elgin, C. (2010). Persistent Disagreement. In Feldman and Warfield (2010), 53–68.
- Feldman, R. (2006) Epistemological Puzzles about Disagreement. In Stephen Hetherington S. (2006). *Epistemology Futures*. Oxford: Oxford University Press, 218–227.
- Feldman, R. (2007). Reasonable Religious Disagreements. In Antony, L. (ed.) (2007). *Philosophers Without Gods: Meditations on Atheism and the Secular Life*. New York: Oxford University Press, 194–214.¹
- Feldman, R. and Warfield, T. (2010). *Disagreement*. Oxford, Oxford University Press.
- Ferri, C. P., Prince, M., Brayne, C., Brodaty, H., Fratiglioni, L., Ganguli, M., Sczufca, M. (2006). Global prevalence of dementia: a Delphi consensus study. *The Lancet*, 366(9503), 2112–2117.
- Goldman, A. (2010a). Social Epistemology. In Zalta, E. (ed.) *The Stanford Encyclopedia of Philosophy*, (online at: <http://plato.stanford.edu/archives/sum2010/entries/epistemology-social/> , accessed April 14, 2015).
- Goldman, A. (2010b). Epistemic Relativism and Reasonable Disagreement. In Feldman and Warfield (2010), 213–26.
- Goldman, A.I. (2001). Experts: Which ones should you trust? *Philosophy and Phenomenological Research*, 63, 85–110.
- Gupta, U. G., & Clarke, R. E. (1996). Theory and applications of the Delphi technique: A bibliography (1975–1994). *Technological forecasting and social change*, 53(2), 185–211.

- James, W. 1896 (1979). The will to believe. In F. Burkhardt et al. (eds.), *The will to believe and other essays in popula philosophy*, Cambridge: MA, Harvard, 291–341.
- Kelly, T. (2005). The Epistemic Significance of Disagreement. *Oxford Studies in Epistemology*, 1, 167-196
- Kornblith, H. (2010). Belief in the Face of Controversy. In Feldman and Warfield (2010), 29-52.
- Kuhn, T. S. (1970). *The structure of scientific revolutions*. Chicago/London: University of Chicago Press.
- Lackey, J. (2008). A justificationist view of disagreement's epistemic significance. *Proceedings of the XXII World Congress of Philosophy*, 53, 145-154.
- Lackey, J. (2010). What should we do when we disagree? In T.S. Gendler & J. Hawthorne (Eds.), *Oxford Studies in Epistemology* (Vol. 3, pp. 274–293). Oxford: Oxford University Press.
- Linstone, H., M. Turoff, M. (1975) *The Delphi Method: Techniques and Applications*. London: Addison-Wesley.
- Miller, B. (2014). Scientific consensus and expert testimony in courts: lessons from the Bendectin litigation. *Foundations of Science*, 1-19.
- Powell, C. (2003). The Delphi technique: myths and realities. *Journal of advanced nursing*, 41(4), 376-382.
- Solomon, M. (2007). The social epistemology of NIH consensus conferences. In H. Kincaid & J. McKittrick (Eds.) (2007). *Establishing medical reality: Essays in the metaphysics and epistemology of biomedical science*. Dordrecht: Springer, 167–177.
- Steup, M., (2010). Epistemology. in: Zalta, E. (Ed.), *The Stanford Encyclopedia of Philosophy*. (online at: <http://plato.stanford.edu/entries/epistemology/>, accessed April 14, 2015).
- van Inwagen, P. (1996). It Is Wrong, Everywhere, Always, for Anyone, to Believe Anything upon Insufficient Evidence. In Jeff Jordan & Daniel Howard-Snyder (eds.), *Faith, Freedom and Rationality*. Savage, Maryland: Rowman and Littlefield (1996), 137-154.

- Wedgwood, R. (2010). The moral evil demons. In Feldman and Warfield (2010), 216-246.
- White, R. (2005). Epistemic Permissiveness. *Philosophical Perspectives* 19/1, 445–59.
- Williams, B. (1973). Deciding to Believe. In *Problems of the Self*. Cambridge: Cambridge University Press.
- Wray, K. (2001). Collective belief and acceptance. *Synthese*, 129(3), 319–333.

