Throwing The Baby Out With The Water: From Reasonably Scrutinizing Authorities To Rampant Scepticism About Expertise

MARKUS SEIDEL

University of Muenster,
Centre for Philosophy of Science,
Domplatz 6, 48143
Muenster, Germany.
maseidel@hotmail.com

Abstract: In this paper, I argue that many arguments from expert opinion are strong arguments. Therefore, in many cases it is rational to rely on experts since in many cases the fact that an expert says that p makes it highly likely that p is true. I will defend this claim by providing five arguments that illuminate and elaborate on five crucial claims about expertise. In this way, I aim to undermine recent attempts to establish a rampant scepticism about arguments from expert opinion.

Résumé: Dans cet article je soutiens que plusieurs arguments fondés sur l'opinion des experts sont des arguments solides. Par conséquent, dans plusieurs cas il est rationnel de s'appuyer sur leurs opinions car, dans de nombreux cas, le fait qu'un expert affirme que *p* rend très probable que *p* est vrai. Je vais défendre l'idée que leurs arguments sont solides en avançant cinq arguments qui éclairent et élaborent sur cinq jugements cruciaux sur l'expertise. De cette façon, je cherche à saper les tentatives récentes d'établir un scepticisme déchaîné sur les arguments fondés sur l'opinion d'experts.

Keywords: arguments from authority, expertise, expert opinion

1. Introduction

In a recent paper in *Informal Logic* Moti Mizrahi has argued that arguments from expert opinion are weak arguments. Mizrahi maintains that empirical evidence on expertise gives us reason to suppose that the fact that an expert says that p does not make it significantly more likely that p is true.

In this paper I take Mizrahi's argument as a vantage point to defend the claim that many arguments from expert opinion are strong arguments. I will provide five arguments that give us reason to suppose that the case for maintaining arguments from expert opinion to be weak is seriously flawed. I will argue that reasonably scrutinizing authorities should not lead us to a

rampant scepticism about expertise by defending the following claims:

a) The property of being an expert should not be confused with the property of being taken to be an expert.

b) The statements made by experts differ as regards the initial probability of their being true.

c) In order to evaluate the reliability of expert opinion it is necessary to distinguish between expertise in different domains.

d) The property of being an expert is a relational property.

e) Rampant scepticism about expertise is self-undermining.

In this way, I will show that the task of philosophical reflection on arguments from expert opinion is to provide criteria to distinguish between weak cases of arguments of expert opinion and those cases in which reference to authority provides inductively strong support – arguments from authority are not fallacious as such.[1]

2. The alleged weakness of arguments from expert opinion

An argument from expert opinion has, according to Mizrahi, the following form:

"(1) Expert E says that p.
 (2) Therefore, p." (Mizrahi 2013, p. 61)

Before I go on to explain what exactly a weak argument is supposed to be and how Mizrahi argues for his conclusion that arguments from expert opinion are weak arguments, let me make a short digression. I think that Mizrahi's formulation is unfortunate for his own purposes.[2] To my mind, a better formulation is the following argument I

---

[1] See also Coleman's paper that appeared in this journal and that is not discussed by Mizrahi (Coleman 1995).

[2] Note, that there is reason to think that Mizrahi's formulation is not just unfortunate *for his own purposes* but *simply* unfortunate. The reason is that even if his simple one-premise-scheme has some initial explanatory value for introducing students to the structure of arguments from expert opinion, it is surely not adequate as a charitable reconstruction of the structure as discussed by his opponents. Thus, for example, Walton and Gordon have argued that the argumentation scheme needs several explicit and implicit premises (assumptions and exceptions) (see Walton and Gordon 2011, p. 11). In this way, they aim to represent different critical questions on specific instances of arguments from expertise as premises in the argumentation scheme by use of their Carneades Argumentation System (http://carneades.github.io/).

(1^I) Expert E says that p.
Therefore,
(2^I) p.

This formulation has the advantage of clearly distinguishing between the question of the *truth* of the premises and conclusion and the *validity* of the argument, which is the question at issue in Mizrahi's paper. According to my formulation, it is possible that even if (1^I) and (2^I) are true the inference nevertheless is a weak one.[3] I am not sure whether this is possible in Mizrahi's formulation because I am not sure what exactly it would mean to say that (2) is true. In any case, it seems to me that Mizrahi also has my inference in mind when he asks his initial question: "is the inference from 'Expert *E* says that *p*' to '*p*' strong?" (Mizrahi 2013, p. 62). Therefore, from now on I will use my formulation of the inference.

It is obvious that those who argue that the inference from "Expert E says that p" to "p" is strong do not intend to claim that the inference is valid *deductively*. The reason is that we regard experts to be *fallible* epistemic sources such that it is possible that though the premises of the inference are true the conclusion is false nevertheless. Take the following deductively valid argument II:[4]

(1^{II}) Everything X says is true.
(2^{II}) X said that p.
Therefore,
(3^{II}) p.

Now, as we know from our encounter with experts, the argument from expert opinion is no instance of this deductively valid argument since we deny that everything an expert says is true; i.e. we assume that experts are fallible epistemic resources. Therefore, the idea behind the claim that the argument from expert opinion is strong is the idea that the fact that an expert

---

I would like to thank an anonymous referee for *Informal Logic* for making me aware of this recent work by Walton and Gordon.

[3] Here is an example: The argument

(1) In the Müller-Lyer-figure it seems to me that the lines are of different length.

Therefore,

(2) In the Müller-Lyer-figure the lines are of equal length,

is a weak one – though it is true that the lines seem to me to be of different length and it is true that the lines are of equal length.

[4] See for the following Walton 1997, p. 92.

says that p *supports* or *makes it more likely that* p (see Mizrahi 2013, p. 61f). This is exactly what Mizrahi denies: "I will argue that premises of arguments from expert opinion do not make their conclusions significantly more likely to be true than false" (Mizrahi 2013, p. 62). If that is true, then arguments from expert opinion are weak arguments, since, according to Mizrahi, "[arguments] from expert opinion are weak arguments unless the fact that expert E says that p makes it significantly more likely that p is true" (Mizrahi 2013, p. 58). I have no objection to this latter claim since it just is a formulation of the close connection between the reliability and the epistemic trustworthiness of an epistemic source. Mizrahi quite correctly points out: "Would you trust a watch that gets the time right 55% of the time? Would you trust a thermometer that gets the temperature right 55% of the time? I suspect the answer to these questions is 'no.' [sic] Similarly, a method of reasoning, such as appealing to expert opinion is trustworthy only if expert opinion is significantly more likely to be true." (Mizrahi 2013, p. 65). The nub of Mizrahi's argument obviously is whether reference to expert opinion is epistemically just as little promising as reference to the mentioned watch or thermometer. Mizrahi thinks so: "expert opinions are only slightly more accurate than chance" (Mizrahi 2013, p. 58). In order to sustain this claim, Mizrahi adduces empirical evidence from the research on expertise. Thus, he claims that "[there] is a growing body of research which shows that experts are wrong more often than one might expect." (Mizrahi 2013, p. 63).

In a nutshell, here is Mizrahi's argument: "I will argue that we do argue fallaciously when we argue that p on the ground that an expert says that p. In other words, I will argue that arguments from expert opinion, i.e., inferences from "Expert *E* says that *p*" to "*p*," [sic] where the truth value of *p* is unknown, are weak arguments. A weak argument is an argument in which the premises, even if true, provide weak support—or no support at all—for the conclusion. Such arguments from expert opinion are weak arguments because the fact that an expert says that *p* does not make *p* significantly more likely to be true. As research on expertise shows, expert opinions are only slightly more accurate than chance and much less accurate than decision procedures." (Mizrahi 2013, p. 58).[5]

---

[5] In what follows, I will only comment on the former part of the last claim from the quote; i.e. only discuss whether expert opinions are only slightly more accurate than chance. My argument will be silent about comparing the reliability of expert opinions to the reliability of decision procedures.

I will argue the contrary: We do not argue fallaciously when we argue that p on the ground that an expert says that p. I will argue that many inferences from "Expert E says that p" to "p", where the truth value of p is unknown to the person making the inference, are strong arguments. This is not to assume that no *argumentum ad verecundiam* is fallacious – a decisive task in informal logic is just to figure out what exactly makes a case of the argument fallacious and another case reasonable (see Walton 1989a, p. 60). The claim that the inferences from "Expert E says that p" to "p" is *as such* a claim that no good reasoning is wrong. Thus, I deny that "the fact that an expert says that p does not make p significantly more likely to be true" (Mizrahi 2013, p. 58).[6] I will provide 5 arguments that will show why we should not throw out the baby of epistemic trust on expertise with the water of reasonably doubting authorities.

3. Argument 1: Arguments from expert opinions are strong arguments because of what it means to be an expert

Take the following characterization of objective expertise: "Person A is an authority in subject S if and only if A knows more propositions in S, or has a higher degree of knowledge of propositions in S, than almost anybody else." (Goldman 1999, p. 268). In this sense, expertise is spelled out in terms of possession of true beliefs: "[...] I first propose that cognitive expertise be defined in 'veritistic' (truth-linked) terms: […] experts in a given domain […] have more beliefs […] in true propositions and/or fewer beliefs in false propositions within that domain than most people do […]" (Goldman 2001, p. 91).[7]

---

[6] I just adopt Mizrahi's phrase here, although I think it is unfortunate: it might wrongly suggest that the proponent of arguments from expertise thinks that saying that p can make p likely true. To my mind a better formulation is that the proponent thinks that the fact that an expert asserts p makes it more likely that p is true; a formulation sometimes used also by Mizrahi (see Mizrahi 2013, p. 58). I will not dwell on that.

[7] Note that Goldman in the context of this quote also uses a veritistic conception of knowledge claiming that knowledge is to be identified with true belief (see Goldman 2009, p. 275). Note also that according to this definition knowledge of more propositions is a necessary and sufficient condition for authority/expertise. At other places, Goldman claims that possession of knowledge is a necessary but not sufficient condition for expertise: in addition, expertise includes skills and techniques to expand knowledge in the domain of expertise (see: "people who have […] a superior quantity or level of knowledge in some domain and an ability to generate new knowledge in answer to questions within the domain." Goldman 2001, p. 91). Since for my argument in the text, it is sufficient that possession of

Now, if we take this definition of what it is to be an expert seriously and further assume that a layperson's opinions in a given domain are about as accurate as chance, then it is an analytic falsity to claim that "expert opinions are only slightly more accurate than chance" (Mizrahi 2013, p. 58). Assume – like we do for epistemic sources that clearly provide strong support – that for the reports of experts appropriate ceteris paribus conditions apply. Thus, just as we should say that visual perception provides strong support in case we aren't on drugs, we aren't short-sighted, there is enough light etc., we should say that the inference from "Expert E says that p" to "p" is obviously fallacious in case E is a notorious liar, E is on drugs, E is threatened by force to claim that p etc. Obviously, nobody taking the inference to be a strong argument is bound to think that the fact that an expert says that p provides support for p in case the expert is threatened by force to say falsities. Surely, it is difficult to spell out exactly the appropriate conditions for trusting in expertise – however, a) that is a problem for trustworthy epistemic sources in general and b) difficulties to spell out the conditions exactly should not lead us to the conclusion that there are no conditions at all.

Now, if we assume that appropriate conditions apply for the inference from expert opinion to be a candidate for providing strong support, then it seems that from Goldman's definition of expertise *it follows* that expert opinions in the expert's domain are far more accurate than chance. To say that expert opinions are only slightly more accurate than chance can be said to mean the following: The probability that p in case expert E says that p and the probability that p in case a layperson says that p after flipping a coin are approximately the same. This fits well with Mizrahi's own characterization: "To say that experts were only slightly more accurate than chance is to say that they might as well have been guessing." (Mizrahi 2013, p. 64). But if that is true with respect to any proposition in a given domain, then we should conclude that E simply is not an expert in that domain. The reason is that an expert in a given domain is someone possessing an extensive fund of true belief and in case the expert is honest, she is sober etc. – i.e. in case appropriate conditions apply – it follows that it is much more probable that p if the expert says that p than if we flip a coin. The argument is simple: in case we find out that certain subjects are as reliable in a specific domain as "a dart-throwing

---

more knowledge is a necessary condition for expertise, I will not dwell on the issue (see Scholz 2009 for discussion).

chimpanzee" (Mizrahi 2013, p. 64), we should conclude that this very fact sustains the conclusion that these subjects are not experts. The reason is that the notion of expertise works in a way that it should not be applied to subjects that are as reliable as dart-throwing chimpanzees in a specific domain. As Paul Woodruff states: "To be an expert is to be someone on whom others may reasonably rely in difficult, perilous or highly technical matters" (Woodruff 1990, p. 68).

However, it is difficult to finally assess the strength of this argument and there are two reasons to be cautious. One reason is that it is much debated exactly what it means to be an expert. Just to focus on Goldman's definition: as Oliver Scholz has shown, Goldman's attempt of a purely veritistic conception of expertise is materially inadequate for the reason that usually a layperson has only few beliefs about the domain at issue whereas the expert has many sophisticated ones. Thus, the expert runs a higher risk of entertaining false beliefs in the domain at issue such that the layperson has at least fewer false beliefs in that domain than the expert (see Scholz 2009, p. 193).[8] Scholz, therefore, proposes that an adequate conception of expertise should not just take into account the dimension of truth but "all epistemic values and desiderata" (Scholz 2009, p. 193), like justification, coherence and understanding.[9] Therefore, since there is no generally accepted account of what it means to be an expert, we should be careful to entertain an argument that draws on assumptions of what it means to be an expert.

A second reason is that the argument just outlined in fact advances some kind of immunization strategy. To see that consider the following exercise in bad philosophy of science: assume a philosopher advances a specific distinction between "scientific" and "unscientific" along a set of criteria C. Assume further that empirical studies in the history and sociology of science conducted in physics laboratories appear to show that the distinction is flawed; assume, for example, that those people we call physicists in part execute their research in a way that is unscientific according to the philosopher's set of criteria C. The historians and sociologists consequently attack the philosopher's distinction and set of criteria C as entertaining unrealistically high demands. If the philosopher now counters this conclusion by claiming that the very fact that the outcome of the empirical studies is that the examined 'physicists' work is unscientific according to criteria C shows that the 'physicists' are not

---

[8] I will say more about this in the next section.
[9] In reply to Scholz, Goldman asserts that he basically accepts this criticism (see Goldman 2009, p. 275).

scientists, the philosopher should be accused of pursuing an immunity strategy. Analogously, it seems, that if we claim that the very outcome of the empirical research advanced by Mizrahi shows that the subjects under study are not really experts, we immunize our philosophical investigation of expertise against challenges from empirical research on expertise.[10] Thus, without scrutinizing the empirical results we cannot assess their impact on the discussion about expertise.

Nevertheless, what the argument emphasizes is that we should distinguish between a subjective and an objective notion of expertise. To my mind, we need to distinguish between somebody being *taken to be* an expert and somebody *being* an expert.[11] This distinction is decisive in order to assess the force of Mizrahi's argument. Mizrahi wants to show that the inference from "Expert E says that p" to "p" is only weak, where a "weak argument is an argument in which the premises, even if true, provide weak support – or no support at all – for the conclusion" (Mizrahi 2013, p. 58). In order to sustain his thesis, Mizrahi provides the example of the FBI estimating that US businesses lose $200-$250 billion to counterfeiting on an annual basis. He claims that the argument from the FBI's saying that US businesses lose that amount of money to the conclusion that US businesses in fact lose that amount is a weak argument because "the FBI 'has no record of source data or methodology for generating the estimate [which] cannot be corroborated.'" (Mizrahi 2013, p. 63). But note that what is shown with this example is just that from the fact that somebody (even everybody) *takes* the FBI to be an authority in estimating the US businesses' loss does not imply that the FBI *is* an authority in this respect. Therefore, Mizrahi's example sustains the obvious fact that the argument from "S is taken to be an expert by another person" to "S is an expert" is fallacious. However, this does nothing to undermine the strength of the inference from "Expert E says that p" to "p": the obvious fact that people can wrongly be regarded to be experts does not sustain the grandiose claim that "expert opinions are only slightly more accurate than chance" (Mizrahi 2013, p. 58).

---

[10] Note that this argument leads on to much debated ground in meta-philosophy; namely the question in which way empirical/scientific research can inspire/influence/falsify philosophical hypotheses. Though I am a proponent of moderate naturalism (see Koppelberg 1996), the argument against the immunization of philosophy in the text should be – and is (see Keil/Schnädelbach 2000) – accepted also by anti-naturalists.
[11] See also Scholz 2009, p. 190f.

4. Argument 2: Experts say many things – some more risky, some quite safe

Let us have a more detailed look on the inference from "Expert E says that p" to "p". Stated in this general form there is a point in saying that the inference is weak, but the reason is rather trivial. It is very plausible that a single person is a layperson in many domains but an expert, if at all, in far fewer domains.[12] Now, obviously the inference from "Expert E says that p" to "p" is weak in case p expresses a proposition in a domain in which E is a layperson: the fact that a person being a complete layperson in estimating outcomes of football games but being a physics Nobel Prize Winner thinks that Bayern Munich will win the Champions League is no reason to think that Bayern Munich will win the Champions League. In short: obviously, it is a requirement for the inference from "Expert E says that p" to "p" to be a candidate for a strong argument that p is a proposition in the field of expertise of E.[13]

Therefore, the argument at issue is the following argument III:

(1^{III}) Expert E in domain D says that p.
(2^{III}) p is a proposition belonging to D.
Therefore,

---

[12] Of course, the truth of this claim depends on how to count domains. Whatever the criterion for distinguishing between domains might be: too broad a criterion risks undermining the distinction between experts and intellectuals (whereas expertise is restricted to a domain, the competence of intellectuals might be said to provide "a contribution to questions that transcend their speciality" (Carrier 2007, p. 14)), too narrow a criterion will have the unwelcome consequence of eliminating the possibility of disagreement among experts because of conceptual reasons.

[13] In fact, this way to misuse arguments from authority is part of the philosophical basis behind the recently much discussed study *Merchants of Doubt* by the historians of science Naomi Oreskes and Erik M. Conway (this book is found also in the references to Mizrahi's text but I have found no reference to the study in the text): the scientists – all physicists – spreading dissent on anthropogenic climate change, the dangers of tobacco smoke and other issues "had no particular expertise in environmental or health questions [but] used their scientific credentials to present themselves as authorities […]. […] Over the course of more than twenty years, these men did almost no scientific research on any of the issues on which they weighed in. […] Yet, for years the press quoted these men as experts, and politicians listened to them, using their claims as justification for inaction." (Oreskes and Conway 2010, p. 8). See for a discussion of the work of Oreskes and Conway in relation to the literature on expertise in philosophy: See Almassi 2012.

(3^{III}) p.

Once stated in this form we can see more clearly what goes wrong in Mizrahi's FBI-case. As far as I can see, this case is an instance of "one common type of error in appeals to expertise in argument. If the expert's field is A, but the issue he is cited as pronouncing upon is another field B, then the argument from authority should be questioned" (Walton 1989b, p. 179). The FBI's expert field is, say, *investigating* and *detecting* counterfeiting, but that should not lead us to expect that the FBI's expert field is *estimating* the US businesses loss to counterfeiting. Probably, many *take* the FBI's expert field to be also estimating the amount lost to counterfeiting, but – as the study cited by Mizrahi shows – they are wrong since the FBI "has no record of source data or methodology for generating the estimate" (Mizrahi 2013, p. 63). We should, therefore, just deny that in the FBI-case premise (2^{III}) is true and should not take this case as an example to sustain that arguments from expert opinion are fallacious.

However, also the restricted argument is not satisfactory to finally assess whether appeal to expert opinion is about as accurate as chance. The reason is that experts say many different things with respect to their domain of expertise and it is of eminent relevance what kind of proposition they assert with respect to how likely it is that the proposition is true. For example, a "successful prediction is valuable because it goes beyond what we already know most directly, but this same feature renders prediction inherently risky." (see Barrett and Stanford 2006, p. 586). Therefore, the case in which an expert says that p, where p is a prediction, and the case in which an expert says that p, where p is a report of an established fact in the field of expertise, differ considerably with respect to the likelihood that p is true. The difference is most obvious in specific domains of expertise; history for example.[14] As far as I see, it is obvious that the following argument is fallacious, therefore: if it turns out that in a specific field the *predictions* of experts are only about as accurate as chance, we have reason to think that the fact that an expert asserts *any statement* p does not make p significantly more likely to be true. But this is exactly how Mizrahi uses some of the evidence from research on expertise: "Indeed, many studies on expertise suggest that the fact that an expert asserts *p* does not make *p* significantly more

---

[14] I will say something about differences in predictions in different domains of expertise in the next section.

likely to be true. For example, Philip Tetlock (2005) conducted a long-term study of numerous *predictions* made by experts from various fields, including academics, economists, policymakers, and journalists. He found that the experts were only slightly more accurate than chance […]." (Mizrahi 2013, p. 64, my italics).[15] Let us assume that Tetlock in fact has established the claim that in the fields studied experts really are no better off in predicting than "a dart-throwing chimpanzee" (Mizrahi 2013, p. 64). Nevertheless, this does not support Mizrahi's claim that the fact that an expert asserts p does not make it significantly more likely that p is true. The reason why this argument is fallacious is simple: assume you have a set of statements E which represents all statements made by somebody in a domain. Assume further that this set can be divided in different sub-sets P and not-P and that the probability of a statement being true in P is lower than in not-P because of some inherent feature shared by all statements in P. However, in that case it is fallacious to argue from empirical data that suggest that the probability of a statement of set P being true is close to 0.5 to the conclusion that the probability of any statement in E being true is close to 0.5. Even shorter: experts do not just make risky predictions.

It might be objected that this argument is a misinterpretation of Mizrahi's claims. The higher risk of predictions being false results from the fact that predictions go beyond what we already know. But, so it might be objected, what Mizrahi means by arguments from expert opinion are just inferences to more risky, unknown areas. Thus, Mizrahi explicitly describes those arguments to be "inferences from 'Expert *E* says that *p*' to '*p*,' [sic] where the truth value of *p* is unknown" (Mizrahi 2013, p. 58). Therefore, the objection is misguided because by "arguments from expert opinion" Mizrahi really means inferences to unknown areas, like e.g. the inference from "Expert E predicts that p" to "p".

In order to assess this objection, we need to ask *to whom* the truth value of p is said to be unknown. If the requirement is that the truth value of p is unknown *to the person making the inference*, then this requirement just is an expression of the conviction that the strength of support of a hypothesis by new data also depends on the degree of belief in the hypothesis before the new data came to light. In our case: for example,

---

[15] See also: "And Yates and Tschirhart […] cite a study by Wilson et al. (1997), which shows that attending physicians caring for elderly patients were no more accurate at *predicting* patients' preferences for end-of-life care than interns […]" (Mizrahi 2013, p. 64, my italics).

expert E's additional invocation of an argument from expert opinion that is meant to sustain "p" by truly affirming "Expert E says that p" is surely a weak argument. The reason is that before invoking the argument expert E obviously knew the truth value of p already; she just reiterates that she knows p and that does not make it any more likely that p is true than before. If Mizrahi wants to restrict the sensible use of the argument from expert opinion to laypersons, then I have no objection[16] – but then, to be sure, the restriction to unknown areas does nothing to restrict his argument just to cases of predictions.

If, however, Mizrahi's restriction to arguments where the truth value of p is unknown is really intended by him to reserve the term "arguments from expert opinion" only for those arguments in which p is a prediction or a proposition hitherto unknown to the expert, then, I think, this restriction misinterprets the statements of the authors he opposes. Thus, for example, Mizrahi quotes Copi, Cohen, and Flage from a logic textbook: "You believe most of what your professors say. When they're speaking within their areas of training and research it is reasonable to do so. They are authorities in their fields […]." (Mizrahi 2013, p. 58). Mizrahi goes on: "In what follows, I will challenge these claims about arguments from authority." (Mizrahi 2013, p. 58). However, I take it that Copi, Cohen, and Flage do not intend to restrict their argument just to predictions by authorities: their reference to the student-professor-situation would be out of place if they did, since usually this situation is characterized by the professor telling facts *known by the scientific community*.

Therefore, my argument stands: since some of the empirical material cited by Mizrahi concerns predictions and since experts do not just make predictions, it is fallacious to think that this material supports the conclusion that arguments from expert opinion are weak.

Note that Mizrahi's apparent focus on predictions by experts has also consequences for his reply to two objections to his conclusion discussed in section 4 of his paper. The second objection discussed there by Mizrahi concerns an argument from

---

[16] Note, that I do not intend to claim that *only* laypersons can make sensible use of the argument from expert opinion. Especially in cases in which expert E believes that p and is confronted with another expert E' believing that not-p and both, E and E', regard the other to be an expert (i.e. a situation that can be called 'revealed peer disagreement' among experts (see Feldman and Warfield 2010, p. 3)), there are difficult questions as to whether because of the very fact of disagreement E and/or E' should split the difference or stick to their guns (see e.g. Elga 2007, Kelly 2010). It goes far beyond the scope of this paper to discuss this topic here.

the history of science: "My overall argument to the effect that arguments from expert opinion are weak arguments flies in the face of the history of science and technology. There has been progress in many disciplines, so the objection goes, and that is reason to think that appealing to expertise is a reliable method of reasoning." (Mizrahi 2013, p. 67). Furthermore, the sixth related objection is that "my argument undermines the scientific enterprise, since science is characterized by a division of cognitive labor" (Mizrahi 2013, p. 74). In reply to this, Mizrahi points out "that scientists themselves rarely, if ever, establish scientific conclusions by appeal to expertise" (Mizrahi 2013, p. 67) and notes that "Einstein didn't argue for the theory of special relativity by appealing to his expertise in theoretical physics." (Mizrahi 2013, p. 67). True as this may be, it misinterprets the force of the original objection. The fact, if it is one, that there is progress in many disciplines is surely partly due to a high degree of division of labour and specialisation in the sciences. Trust in the reliability of experts is a basic mechanism that makes these latter features possible. Furthermore, trust in what experts say is a condition of becoming a scientist at all – if not, why study physics at the university and not by flipping coins – and making progress on your own.[17] This is the rationale behind Newton's aphorism: "If I have seen further it is by standing on the shoulders of giants."[18] Surely, no scientist should establish conclusions just by pointing out that he says so but – as Mizrahi correctly notes – by "appeal to observations and experiments" (Mizrahi 2013, p. 67). But progress in science is also due to trust in the expertise of others that report the results of experimentation and observation – just take the example of the article written by 99 authors that is cited in John Hardwig's seminal essay 'Epistemic Dependence' (see Hardwig 1985, p. 346). Mizrahi unduly focuses on expertise as a means to establish new conclusions or to predict but seems to forget about other kinds of knowledge from expertise. The fact that some kinds of statement by experts are more likely to be false should

---

[17] Note that this argument is not an argument necessarily made by scientific realists. For example, I think it is not difficult to illuminate the notion of progress allowed by Thomas Kuhn in his discussion of normal science along this line. This is reason to doubt that the second reply by Mizrahi to the objection has much force: he notes that the question of whether there is progress in terms of approximation to truth in science is still debated in philosophy of science (see Mizrahi 2013, p. 67). However, if we can reconstruct the argument also in terms of Kuhn's notion of progress, this debate should not bother us too much.
[18] See Merton 1965 for discussion.

not lead to the conclusion that arguments from expert opinion are weak.[19]

5. Argument 3: An account of expertise should take into account the specific nature of the field of expertise

Let me discuss Mizrahi's empirical evidence for his conclusion in more detail. There is one remarkable feature of nearly all studies cited by Mizrahi in order to sustain his thesis: they concern expert-judgment in fields that probably are special with respect to the reliability of their results and predictions. Thus, the examples are from the fields of politics, medical science, psychology and economics (see Mizrahi, sect. 3). It would be clearly beyond the scope of this paper to discuss the differences and similarities between these disciplines and other disciplines like e.g. physics or chemistry. Nevertheless, there are differences especially with respect to the reliability of predictions. Thus, for example, Nobel Prize Winner for economics Amartya Sen notes that "economic predictions are notoriously unreliable" (Sen 1986, p. 3). Sen further explains the difference between predictions in physics and economics by referring to the complexity of the field: "The real difference lies in the enormous variety of 'respects' that could actually make an important difference in economics, and the complex ways in which these economic influences operate and interact with each other" (Sen 1986, p. 4). It is not implausible to say that also in the other fields from which Mizrahi takes the empirical evidence on expertise – namely politics, medical science and psychology – experts are faced with a similar problem of complexity as noted by Sen for the case of economics.[20]

Surely, much more needs to be said in order to assess whether Sen's explanation of the difference between physics and

---

[19] In addition, Mizrahi argues that "scientists, for the most part, don't have to appeal to expertise to justify their trust in the instruments that they use because (a) respectable scientists probably know how to test the instruments that they use" (Mizrahi 2013, p. 67). I simply doubt that. Apart from easy ways of knowing how to test the instrument they are using (e.g. pressing the test-mode-button) I do not think that scientists usually have enough knowledge about the workings of the instruments they are using such that they can test them for themselves; probably in case the instrument seems to work improperly and it thus should be tested, the scientist will call a technician to test the instrument. (I admit, however, that only empirical evidence can finally decide the case).

[20] Mizrahi also notes one example from physics (Mizrahi 2013, p. 63). I will discuss that example in section 7.

economics is right and whether his explanation applies also to the fields of medical science, politics and psychology. Nevertheless, in order to use the fact that experts *in economics, in medical science, in psychology and in politics* are only slightly more accurate than chance as evidence for the conclusion "that *experts* are only slightly more accurate than chance" (Mizrahi 2013, p. 64, my italics), the possibility that the unreliability of experts in these domains stems from specific features of these domains needs to be excluded.[21]

My argument should not suggest that we can only have reliable expertise in domains like the hard sciences. Take the example of football. Although I have not conducted an empirical study concerning the reliability of football experts, I think it is highly plausible that football experts are much more often right in predicting the outcome of football matches than laypersons relying on coin flips. Let us assume that in the Chinese Super League we want to predict the outcome of the match between Guangzhou Evergrande and Henan Construction.[22] Now assume that we have no clue about Chinese football and we are given the choice between flipping a coin to make a prediction and asking Juan Antonio Camacho, who is the recent coach of the Chinese national team,[23] about his opinion, it is surely highly rational to ask Camacho. The reason is not that Camacho cannot be wrong in his prediction – he is surely no clairvoyant – but that he is an expert in Chinese football in comparison with us. In contrast to the coin, Camacho knows that Guangzhou has invested a lot of money to buy high quality players like Lucas Barrios and Dario Conca, whereas Henan mainly recruits players from the Chinese market. Therefore, in contrast to us Camacho can take this knowledge into account by estimating whether Guangzhou or Henan will win the match. And since the just mentioned information about the quality of the players obviously affects the probability that Guangzhou will win the

---

[21] This problem is not restricted to predictions: Mizrahi also notes that most results published in medical and in economics journals are rejected after a few years (see Mizrahi 2013, p. 64). In order to constitute evidence for Mizrahi's thesis it must be assured that these facts are not due to specific features of these disciplines; e.g. by a similar study with respect to physics journals. See for discussion of Mizrahi's argument also the next section.

[22] I choose the Chinese Super League because probably most of the readers will be complete laypersons with respect to the strength of the teams.

[23] Actually, Camacho has been sacked as coach of the team whilst this paper was under review. Of course, this fact does not change anything in the argument.

match, it is highly plausible that Camacho will say that Guangzhou will win the match.[24]

The example just mentioned comes from an area in which there is an enormous variety of factors that could make a difference in predicting the outcome and that also Camacho does not or cannot take into account by making his prediction: the exact length of the lawn, the weather conditions, the exact tactical considerations of Marcello Lippi, who is Guangzhou's coach and so on. Nevertheless, I think it is highly plausible to assume that Camacho is not "only slightly more accurate than chance" in predicting the outcome of matches between Guangzhou and Henan. The inference from "Camacho says that Guangzhou will win the match" to "Guangzhou will win the match" is a good one, although it surely is possible that Henan will win the match.

My objection is not meant to suggest that the empirical results that show that experts are wrong more often than we think are irrelevant, false or less disturbing. On the contrary, checking the track-record of experts by empirical research is of major importance for assessing the reliability of experts – empirical studies on the reliability of assumed experts provide the layperson with decisive evidence on which experts to trust and whether to trust on experts in specific domains.[25] If, to take up again the example from Chinese football, empirical research on Camacho's predictions about the outcome of Chinese football matches shows that Camacho is about as accurate in predicting as chance, then this evidence should lead us to doubt Camacho's expertise. Note, that this is also what is suggested by the empirical studies cited by Mizrahi: thus, for example Freedman 2010 closes his critical book on expert advice by providing general rules of thumb of "more trustworthy expert advice" (Freedman 2010, p. 224). Though I disagree with Freedman about his concrete suggestions, I agree that this is the right focus of debate about expert opinion: How can a novice identify an expert? How can a novice decide between two putative experts?[26] How to discriminate between experts and

---

[24] Note, that in case Camacho honestly says that Henan will win it is rational to conclude that Henan will win *even in case we know the facts about Guangzhou and Henan mentioned in the text*. The reason is that we take Camacho's surprising prediction itself to be evidence for the fact that Camacho has some additional, decisive information unknown to us.

[25] See Scholz 2009, p. 201f.

[26] See for discussion of these questions – the novice/expert and the novice/2-expert problem – Goldman 2001 and Scholz 2009.

frauds?[27] Therefore, the empirical evidence cited by Mizrahi is highly disturbing as regards our trust in expertise in specific domains. It is, however, no reason to conclude quite generally that experts are only slightly more accurate than chance – it rather invokes us to work on providing criteria for identifying experts and for distinguishing between trustworthy and not-trustworthy testimonial acts.

6. Argument 4: Being an expert is a relational property

Whenever I enter my sports club, my comrades greet me with "Here's our philosophy expert". Unfortunately, this never happens when I arrive at a philosophy conference. What is the explanation of this difference given the fact that the amount of true beliefs that I entertain surely does not change once I enter my sports club or a conference? Surely, part of the explanation of the difference in reaction is my comrades' mocking at me and the different communicational situations at the sports club and at a scientific conference. However, part of the explanation is also that my sports comrades are simply right that *in relation to them* I am a philosophy expert whereas *in relation to most of the other participants* on the conference I am at best a peer as regards philosophy. Therefore, it is plausible to assume that whereas somebody can be an expert in domain D in one context the same person can be a non-expert in domain D in another context. The property of being an expert is a *relational property*.

However, that does not mean that in the land of the blind, the one-eyed man is king. Though being an expert is a relational property, expertise is not just context-relative. An account of expertise should also specify that it is a necessary condition for being an expert to exceed a certain minimum of believing true and justified beliefs, of understanding and other epistemic desiderata. Surely, the smartest kid in primary school is not simply an expert by the fact that she outplays her peers.

An exact specification and explication of the property of being an expert would be far beyond the scope of this paper, but I think that the following two conditions are necessary conditions for an account of expertise:
a) Somebody is an expert only *in relation* to a person or group of people.

---

[27] This question is not a new one in philosophical debate: See Gentzler 1995 for discussing this question in relation to Plato's early dialogues.

b) Somebody is an expert only if she *exceeds a minimum* of epistemic desiderata.[28]

Necessary condition b) is necessary in order to meet a requirement on epistemic sources in general: epistemic sources should be reliable guides to the truth. If we drop condition b) such that in the land of the blind the one-eyed man is king, we have no reason to assume that expertise is an epistemic source: though being king of the blind, the one-eyed man still is visually handicapped.

The focus of Mizrahi's argument is mainly on showing that condition b) is not fulfilled in a sufficient number of cases of putative expertise. However, it is important not to forget condition a). And, most importantly, it is important to see that – though independent of each other – there are subtle relations between condition a) and b). Just take the fact that the minimum of true and justified beliefs required to fulfil condition b) partly depends on the specification of the group of people in condition a). This is important in assessing the expertise of scientists in the history of science. It would be not in tune with our concept of expertise to demand of scientists of past epochs that they go far beyond the best beliefs of the scientific community at their time in order to be experts. Thus, to say that Aristotle is not an expert in biology *because* he did not know that DNA has the structure of a double-helix is no good argument. Even in case we finally come to the conclusion that Aristotle cannot be regarded as an expert in biology surely our conclusion cannot be backed up by his lack of knowledge of the structure of DNA. A reasonable criterion of the required amount of epistemic desiderata should take into account the relational structure of expertise.

If that is true, however, we should be careful not to overstate the fact that in most disciplines in science the accepted state of research is evolving continuously. Especially at the very frontline of research it is no surprise that many published results finally turn out to be false or need some revision after some time. This is no shortage of science but part of the very mechanism by which progress in science becomes possible. Therefore, the fact that scientific change demands evaluating the alleged results of past science to be largely incorrect does not *by itself* sustain the claim that those holding the incorrect claims

---

[28] See for example the definitions in Scholz 2009, p. 192f. that fulfill both necessary conditions. Of course it is far from easy to incorporate both necessary conditions into an elaborate account of expertise – there seems to be a subtle tension between these conditions. To my knowledge, a satisfactory theory of expertise that accounts for both intuitions underlying these conditions is still a desideratum.

have not been trustworthy experts. The relational character of expertise demands that we do not just simply count the statements of alleged experts along the different epistemic relevant dimensions like truth, justification and so on, but that we do so in relation to the actual state of research. If that is true, then we should reassess some of the evidence cited by Mizrahi however. Mizrahi notes that research on expertise has shown that in some fields most of the results published in top journals are rejected after a few years (see Mizrahi 2013, p. 64). Does this fact sustain the claim that experts in these fields are only slightly more accurate than chance? Not necessarily, since this fact might also point out to the evolving, and in this way *scientific*, character of the field.[29] Most importantly, the relational character of expertise should be reason enough not just to evaluate expertise by epistemically assessing the beliefs of somebody held at time t in relation to the beliefs held at a later time t', but also to evaluate expertise by epistemically assessing the beliefs of somebody held at time t in relation to the beliefs of others held at time t. Since a) and b), as far as I see, are both necessary conditions of expertise, any argument that aims to show that condition b) is not met but does so by forgetting about a) is not convincing.

7. Argument 5: Radically doubting expertise is self-undermining

In this section, I am going to argue that radical scepticism about expertise is virtually impossible because it is self-undermining. The point is neatly summarized by Evan Selinger and Robert P. Crease in their introduction to the volume *The Philosophy of Expertise*: "Ironically, even if it were desirable to be sceptical about expert pronouncements, such scepticism could never become radical. The ability to doubt particular expert claims necessitates appealing to an alternative base of knowledge, much of which must also be imparted by experts." (Selinger/Crease 2006, p. 2).

---

[29] Note that I do not intend to claim that this is true for the fields mentioned by Mizrahi, namely medicine and economics (see Mizrahi 2013, p. 64). Perhaps reflection and empirical research in these areas finally shows that in these areas we should not trust arguments from expert opinion. However, I just want to point out that the fact that many results published in top journals in a field are rejected after a few years does not imply that experts in these fields are not reliable epistemic sources.

In order to see the problem let us note that the following argument IV surely is a good argument:

(1^{IV}) Evidence gained by epistemic method E' suggests that epistemic method E" is not reliable.
(2^{IV}) Epistemic method E' is a reliable method.
Therefore,
(3^{IV}) There is evidence that suggests that epistemic method E" is not reliable.
Therefore,
(4^{IV}) Epistemic method E" is not reliable.

We make use of arguments of this kind very often in our epistemic practice – this argument is part of bringing our epistemic house in order. However, note that the argument is only a good argument if an implicit condition is met. Thus, the following argument V is not a good argument:

(1^{V}) Evidence gained by epistemic method E' suggests that epistemic method E" is not reliable.
(2^{V}) Epistemic method E' is a reliable method.
(3^{V}) As a matter of fact, the evidence gained by epistemic method E' that suggests that epistemic method E" is not reliable is evidence gained by assuming the reliability of epistemic method E".
Therefore,
(4^{V}) There is evidence that suggests that epistemic method E" is not reliable.
Therefore,
(5^{V}) Epistemic method E" is not reliable.

It is important to be precise about premise (3^{V}) here because there is an argument similar to argument V that is also fallacious but it is so because it is self-contradictory. This is the argument that we arrive at by substituting the following premise for (3^{V}) in argument V:

($3^{V'}$) Epistemic methods E' and E" are identical.

Such a reformulated argument is self-contradictory because it is impossible in this argument that both the premises and the conclusions are true: from (2^{V}) and ($3^{V'}$) it follows that epistemic method E" is a reliable method and this contradicts (5^{V}).

In order to see the difference between argument V and the reformulated version of it let us assume that we aim to show that sensory perception is not a reliable epistemic method. Assume we reason as follows:

(1^{SENS}) Evidence gained by sensory perception suggests that sensory perception is not reliable.
(2^{SENS}) Sensory perception is a reliable epistemic method.
Therefore,
(3^{SENS}) There is evidence gained by a reliable method that suggests that sensory perception is not reliable.
Therefore,
(4^{SENS}) Sensory perception is not reliable.

This argument is fallacious because it is impossible that both the premises and the conclusion of this argument are true – (2^{SENS}) and (4^{SENS}) are contradictory.

Now assume we reason for the same conclusion not along the lines of the foregoing argument but on the following lines:

(1^{TEST}) Evidence gained by testimony suggests that sensory perception is not reliable.[30]
(2^{TEST}) Testimony is a reliable epistemic method.
Therefore,
(3^{TEST}) There is evidence gained by a reliable method that suggests that sensory perception is not reliable.
Therefore,
(4^{TEST}) Sensory perception is not reliable.

This argument is also fallacious but not because it is impossible that both the premises and the conclusion can be true. It is fallacious because it is an instance of the original argument V. For in order to gain testimonial evidence we must assume that at least sometimes sensory perception reliably leads to the truth: in any case, testimonial evidence must be heard or seen. Therefore, the argument is fallacious because it implicitly assumes the falsity of the conclusion; it is not fallacious because it is impossible that both the premises and the conclusion can be true.

On the basis of these considerations, I maintain that the argument for radical scepticism about expertise and Mizrahi's argument are instances of the original argument V. I argue that it

---

[30] Just imagine the case of the movie *The Matrix* where Neo is told by Morpheus that the visible world is an illusion (see Goldman 2010, p. 200 for this example).

is virtually impossible to argue that arguments from expert opinion are weak because such arguments implicitly assume that arguments from expert opinion are strong. In short, it is question-begging to use evidence gained assuming the reliability of an epistemic method in order to sustain the claim that this epistemic method is not reliable.

Let us assume that epistemic method E'' consists in reference to expert opinion in order to justify claims. Let us further assume that epistemic method E' consists in what can be roughly described as 'scientific and empirical research'. Thus, the following is a fallacious argument:

(1^{VI}) Evidence gained by scientific and empirical research suggests that reference to expert opinion is not reliable.
(2^{VI}) Scientific and empirical research is a reliable method.
(3^{VI}) As a matter of fact, the evidence gained by scientific and empirical research that suggests that reference to expert opinion is not reliable is evidence gained by assuming the reliability of reference to expert opinion.
Therefore,
(4^{VI}) There is evidence that suggests that reference to expert opinion is not reliable.
Therefore,
(5^{VI}) Reference to expert opinion is not reliable.

It is obvious that Mizrahi subscribes to premise (1^{VI}) and (2^{VI}).[31] However, is premise (3^{VI}) true in Mizrahi's case? Is Mizrahi assuming that reference to expert opinion is reliable in order to gain the evidence from empirical research? In fact, Mizrahi does. As Mizrahi admits, he has not conducted any experimental studies on expertise: "Luckily, I don't have to. Others have done the hard work already." (Mizrahi 2013, p. 76). Therefore, Mizrahi is relying on the expertise of others in conducting empirical studies on expertise in order to come to his claim that there is empirical evidence for the conclusion that arguments from expert opinion are weak arguments. Thus, Mizrahi in fact refers to expert opinion in order to sustain his premise (1^{VI}) and therefore Mizrahi's argument is an instance of argument VI.

In order to see this let us have a look at one of Mizrahi's examples. I quote in some length:

> In 1989, Dr. Martin Fleischmann and Dr. Stanley Pons, both electrochemists working at the University of Utah at

---

[31] See e.g. Mizrahi 2013, p. 58.

the time, announced that they had found a way to create nuclear fusion at room temperature. Suppose, then, that, shortly after their announcement, a non-expert puts forward the following argument from expert opinion:

(1) Electrochemists Fleischmann and Pons say that nuclear fusion can occur at room temperature.

(2) Therefore, nuclear fusion can occur at room temperature.

In this case, a true premise in an argument from expert opinion leads to a false conclusion. For, as it turns out, nuclear fusion cannot occur at room temperature. […] As it turned out, when other experts tried to replicate the results reported by Fleischmann and Pons, they could not do so. (Mizrahi 2013, p. 62).

Note how Mizrahi argues for the falsity of premise (2): because other *experts* failed to replicate the results, Mizrahi believes that nuclear fusion cannot occur at room temperature.[32] How does he know? Probably because the experts said that they could not replicate the results. Mizrahi, therefore, advances the following argument VII:

(1^{VII}) Experts say that they could not replicate the results of Fleischmann and Pons.
Therefore,
(2^{VII}) The results of Fleischmann and Pons are not replicable by experts.
(3^{VII}) Non-replicability of results by experts is evidence for the falsity of the results.
Therefore,

---

[32] It might be objected that my reading of the quoted passage is uncharitable: Mizrahi might not mean his reference to experts to serve an epistemic function such that he *justifies* his belief in the non-replicability of the results by reference to experts but that he just aims to make a claim about the *genesis* of his belief and that the sole epistemic role for his argument is played by the actual non-replicability (this, in fact, takes up Mizrahi's own defense against the reproach of self-contradiction, see Mizrahi 2013, p. 76). However, I do not think that such a weaker reading is suggested here as will be clear if we try to reconstruct Mizrahi's argument more coherently. If his reference to experts does not play a justificatory role here, then the plausibility of his argument should not be reduced by substituting "randomly chosen people" for "experts". But I think it is obvious that the following argument is far from plausible: "As it turns out, nuclear fusion cannot occur at room temperature. As it turned out, when randomly chosen people tried to replicate the results reported by Fleischmann and Pons, they could not do so.". Therefore, I believe that my reading is not uncharitable – Mizrahi in fact trusts that the other experts' inability to replicate *justifies* the belief that the results are not replicable.

(4^{VII}) Nuclear fusion cannot occur at room temperature.

The step from (1^{VII}) to (2^{VII}) is an argument from expert opinion used by Mizrahi. If Mizrahi really believes in his conclusion that arguments from expert opinion are weak arguments, then it is miraculous how he aims to sustain the step from (1^{VII}) to (2^{VII}).[33]

It is important to notice that Mizrahi's fallacious argument VI should not be confused with another fallacious argument. The following argument VIII is fallacious too:

(1^{VIII}) Experts on expertise say that reference to expert opinion is not reliable.
(2^{VIII}) Reference to expert opinion is a reliable method.
Therefore,
(3^{VIII}) There is evidence that suggests that reference to expert opinion is not reliable.
Therefore,
(4^{VIII}) Reference to expert opinion is not reliable.

In his discussion of anticipated objections Mizrahi correctly notes that his argument is not an instance of this fallacious argument: "My argument for the second premise in my overall argument is not "Experts on expertise say that expert opinions are unreliable; therefore, expert opinions are unreliable." Rather, my argument is that *empirical evidence* shows that expert opinions are unreliable." (Mizrahi 2013, p. 76). It is true: Mizrahi does not apply an argument along the lines of the reformulated argument V that includes premise ($3^{V'}$). But nevertheless Mizrahi's argument is self-undermining because in order to gain the empirical evidence that allegedly shows that expert opinions are unreliable Mizrahi assumes that

---

[33] It might be objected that the foregoing argument differs in kind from the argument using testimonial knowledge. The reason is that in the case of testimonial knowledge it is impossible to get evidence from testimonial sources without assuming the reliability of sensory perception whereas the fact that Mizrahi uses expert knowledge does not imply that it is impossible to compile scientific evidence without relying on arguments from expert opinion.
The objection is right and the difference is important. Nevertheless, note that my claim is that it is *virtually* impossible to be a radical sceptic about expertise. In particular, I do not see how *scientific* research can sustain such a radical skepticism: as Newton says, the actual scientist can see further by standing on the shoulder of giants. In this way, it is surely not impossible *per se* to see further without standing on the shoulder of giants but it is *virtually* impossible.

arguments from expert opinion are strong arguments. He assumes that the following argument from expert opinion is a strong argument; i.e. its premise makes it more likely that the conclusion is true:

(1^{EXP}) Experts different from Fleischmann and Pons say that they could not replicate the results by Fleischmann and Pons.
Therefore,
(2^{EXP}) The results by Fleischmann and Pons could not be replicated by other experts.

I do not see how to evade Selinger's and Crease's argument cited at the beginning of this section: entertaining empirical evidence on expertise actually presupposes the general reliability of experts such that using the empirical evidence to sustain the claim that experts are as reliable as flipping coins is self-undermining.

8. Conclusion

In this paper I elaborated on the question whether arguments from expert opinion are weak arguments. I provided 5 arguments that speak against such a general conclusion. These arguments are backed up by 5 claims about expertise:
a) The property of being an expert should not be confused with the property of being taken to be an expert.
b) The statements made by experts differ as regards the initial probability of their being true.
c) In order to evaluate the reliability of expert opinion it is necessary to distinguish between expertise in different domains.
d) The property of being an expert is a relational property.
e) Rampant scepticism about expertise is self-undermining.
A sensible philosophical treatment of the structure of arguments from expert opinion and the relevant epistemological questions concerning testimony from experts should flinch from drawing rampant sceptical conclusions from epistemological problems in the philosophy of expertise. The task of evaluating the reliability of epistemic sources consists partly in specifying criteria of adequacy for the application of the source. The empirical investigation of expertise is necessary in order to find out which authorities and experts to trust. Justified scepticism concerning some 'experts' should however not be used to justify rampant scepticism about expertise.

Acknowledgements

References

Almassi, B. (2012). "Climate Change, Epistemic Trust, and Expert Trustworthiness." *Ethics & the Environment* 17 (2): 29-49.

Barrett, J. A. and P. K. Stanford (2006). "Prediction." In S. Sarkar and J. Pfeifer (Eds.), *The Philosophy of Science. An Encyclopedia.* (pp. 585-599), New York/London: Routledge.

Carrier, M. (2007). "Engagement und Expertise: Die Intellektuellen im Umbruch." In M. Carrier and J. Roggenhofer (Eds.): *Wandel oder Niedergang? Die Rolle der Intellektuellen in der Wissensgesellschaft.* (pp. 13-32), Bielefeld: transcript.

Coleman, E. (1995). "There is no Fallacy of Arguing from Authority." *Informal Logic* 17 (3): 365-383.

Elga, A. (2007). "Reflection and Disagreement." *Nous* 41 (3): 478-502.

Feldman, R. and T. A. Warfield (2010). "Introduction." In R. Feldman and T.A. Warfield (Eds.): *Disagreement*. (pp. 1-9), New York: Oxford University Press.

Freedman, D. H. (2010). *Wrong. Why Experts Keep Failing Us – and How to Know When not to Trust Them.* New York: Little, Brown and Company.

Gentzler, J. (1995). "How to Discriminate Between Experts and Frauds: Some Problems for Socratic Peirastic." *History of Philosophy Quarterly* 12 (3): 227-246.

Goldman, A. (1999). *Knowledge in a Social World.* Oxford: Oxford University Press.

Goldman, A. (2001). "Experts: Which Ones Should You Trust?" *Philosophy and Phenomenological Research* 63 (1): 85-110.

Goldman, A. (2009). "Replies to Discussants." *Grazer Philosophische Studien* 79, 245-288.

Goldman, A. (2010). "Epistemic Relativism and Reasonable Disagreement." In R. Feldman and T.A. Warfield (Eds.): *Disagreement.* (pp. 187-215), New York: Oxford University Press.

Hardwig, J. (1985). "Epistemic Dependence." *The Journal of Philosophy* 82 (7): 335-349.

Keil, G. and H. Schnädelbach (2000). Naturalismus. In G. Keil and H. Schnädelbach (Eds.): *Naturalismus. Philosophische Beiträge.* (pp. 7-45). Frankfurt: Suhrkamp.

Kelly, T. (2010). "Peer Disagreement and Higher-Order Evidence." In R. Feldman and T.A. Warfield (Eds.): *Disagreement.* (pp. 111-174), New York: Oxford University Press.

Koppelberg, D. (1996). "Was Macht eine Erkenntnistheorie Naturalistisch?" *Journal for General Philosophy of Science* 27 (1): 71-90.

Merton, R. K. (1965). *On the Shoulders of Giants. A Shandean Postscript.* New York: The Free Press.

Mizrahi, M. (2013). "Why Arguments From Expert Opinion are Weak Arguments." *Informal Logic* 33 (1): 57-79.

Oreskes, N. and E.M. Conway (2010). *Merchants of Doubt. How a Handful of Scientists Obscured the Truth from Tobacco Smoke to Global Warming.* New York: Bloomsbury Press.

Scholz, O.R. (2009). "Experts: What They Are and How to Recognize Them – A Discussion of Goldman's Views." *Grazer Philosophische Studien* 79, 187-205.

Selinger, E. and R.P. Crease (2006). "Introduction." In E. Selinger and R.P. Crease (Eds.): *The Philosophy of Expertise.* (pp. 1-7), New York: Columbia University Press.

Sen, A. K. (1986). "Prediction and Economic Theory." *Proceedings of the Royal Society of London A*, 3-23.

Walton, D. N. (1989a). "Reasoned Use of Expertise in Argumentation." *Argumentation* 3, 59-73.

Walton, D. N. (1989b). *Informal Logic. A Handbook for Critical Argumentation.* New York: Cambridge University Press.

Walton, D. N. (1997). *Appeals to Expert Opinion. Arguments from Authority.* University Park: Pennsylvania State University Press.

Walton, D. N. and T. F. Gordon (2011). "Modeling Critical Questions as Additional Premises." In F. Zenker (Ed.): *Argument Cultures: Proceedings of the 8th International Conference of the Ontario Society for the Study of Argumentation (OSSA).* (pp. 1-13), download at: http://www.dougwalton.ca/papers%20in%20pdf/11OSSA.pdf

Woodruff, P. (1990). "Plato's Early Theory of Knowledge." In S. Everson (Ed.): *Companions to Ancient Thought: Epistemology.* (pp. 60-84), Cambridge: Cambridge University Press.