

*Martin Hoffmann*

## How to Identify Moral Experts? An Application of Goldman's Criteria for Expert Identification to the Domain of Morality

*Abstract:* How can laypeople justifiably distinguish between reliable experts and unreliable experts? This problem, usually called the 'problem of expert identification', is highly debated in recent social epistemology. A great amount of work has been undertaken in order to find satisfactory criteria for identifying experts in different branches of the empirical sciences, but hardly in the domain of moral knowledge. This asymmetry between social and moral epistemology is the motivation behind my paper. I reconsider the epistemological problem of identifying moral experts by applying identification criteria developed in general social epistemology to the area of morality. As I will show, all of these criteria turn out to be inappropriate for identifying *moral* experts. This result seems implausible, because it conflicts with the observation that moral experts play an important role in public and scientific discourse, in ethics committees and boards. But this is not a real contradiction—as I will illustrate by explaining which tasks these experts can, in my view, fulfil.

### 1. Introduction

The German novelist and essayist Ingo Schulze wrote an essay about the financial crisis in the European Union. He reported two observations he made in panel discussions on this controversial topic with eminent economical, political and moral implications: first, the impact or authority of statements in the course of these discussions hinges essentially on the (economical) expertise ascribed to the speaker by the audience. Second, every time a person—often Schulze himself, in this case—expresses a statement which is not in line with the established opinions (which are usually shaped by the economical experts), these statements are primarily not taken as a dissenting opinion on the matter, but as an information about the person who made the statement. This 'switching' effect (as I will call it) has the consequence that the dissenting statement is not interpreted as a contribution to the topic in question that is worth considering, but as an expression of the worldview (naïve, socialist or a mixture of both?) of the speaker (Schulze 2012, 11–13, 39–40).

Is this observation of any philosophical relevance? I think it is. It draws our attention to some important aspects of knowledge acquisition in social contexts,

which is the subject of social epistemology. The first aspect is an epistemological truism: the epistemic authority of a statement expressed by a speaker is *inter alia* determined by the status of expertise ascribed to her or him in a certain social context. For this reason the concepts of *expertise* and *expert* are highly debated in recent social epistemology and traditionally the concept of epistemic authority was a starting point for several definitions of the notion of an expert (Scholz 2009, 189–191). The second aspect is a more delicate one and deserves further philosophical clarification: verbal statements of individual speakers not only inform us about facts, but can also contain information about the speaker herself. This phenomenon even occurs in everyday situations. A simple example: imagine me and four other people looking at a red house. We all agree: “The house is red!” A sixth person joins us and states: “The house is grey!” Obviously, we do not learn anything new about the house’s colour, but we are informed about the colour-blindness of the sixth person.

I quote Schulze’s observation of the ‘switching’ effect, because it gives rise to a fundamental epistemological question: when am I justified to take a statement as a reliable testimony and when should I attribute the statement to failures or shortcomings of the speaker? This question may be easily answered in everyday contexts like in my example, but it becomes more complicated in situations where our own knowledge is so limited that we have to rely on expertise. Currently, we need expertise in many areas to solve the pressing problems of our time. If we want to evaluate the consequences of the European financial crisis we need economic expertise on the one hand, but also political expertise in order to be informed about the enforcability of governmental interventions, and also moral expertise in order to consider what is permissible to do and to which actions we might be morally obliged while our economy and currency system is in danger.

The epistemological key question is: how can we decide which experts to trust and which to distrust? In recent literature in social epistemology this problem is discussed under the title the “problem of *expert identifiability*” (Goldman 1991, 129). In general, the theoretical task of solving this problem consists in the development of a catalogue of criteria to differentiate between experts and laypeople in a certain domain. Up to now a great amount of work has been undertaken in order to find satisfactory criteria for identifying experts in different areas of the empirical sciences (Goldman 1999; 2001; Scholz 2009). By contrast, there is comparatively little work done to solve the problem of expert identifiability for the domain of moral knowledge. In the recent literature on the topic moral experts the identification problem is often evaluated as “not a particularly deep or in principle an insurmountable problem” (Hills 2009, 96). The current discussion in moral philosophy locates the main problem not in the epistemic unavailability of moral knowledge by testimony, but in its *unusability* (Hopkins 2007, 626ff.). This tendency leads to a shift away from epistemological questions and to a focus on the moral problem of recognizing the authority of experts in moral questions while remaining an autonomously acting person, claiming one’s right of democratic self-government in modern societies (Archard 2011, 125–127).

My main project in this paper is to reconsider the epistemological problem of identifying moral experts. I apply the most important identification criteria developed in general social epistemology to the area of morality. As I will show, all of these criteria turn out to be inappropriate for identifying *moral* experts. Hence my considerations undermine the credibility of alleged moral experts in general. There are no good epistemic reasons to believe in their moral advice.

At first glance, this result might seem absurd. The appeal to moral expertise is highly established in philosophical ethics and public discourse. It is deeply entrenched in medical and research ethics (clinical ethics committees, institutional review boards), risk assessment in environmental ethics (see e.g. the ethics committee for a safe energy supply, *Ethikkommission für eine sichere Energieversorgung*, installed by Federal Chancellor Angela Merkel on 22 March 2011) and many other areas, especially in applied ethics. So my second task is to explain why this social practice is not necessarily a contradiction to my arguments.

In the following section I will introduce a set of acknowledged criteria developed for the identification of experts in general social epistemology (*section 2*). Then I will apply these criteria to the area of moral expertise and will show that the arguments for the credibility of moral experts do not succeed to justify the trust in moral experts (*section 3 to 5*). I will conclude by explaining which role moral experts in ethics committees and boards may nevertheless fulfill (*section 6*).

## 2. What Is an Expert and How to Identify Her?— Criteria for Expert Identification

In social epistemology the concept of expertise is usually characterized as a special case of testimony: there are many instances of knowledge by testimony which are not instances of knowledge on the basis of expertise, but every instance of knowledge by expertise is an instance of knowledge by testimony. For this reason, an expert is an epistemic witness who fulfils some further conditions. In what follows, I apply, to a great extent, Goldman's definition of and identification criteria for experts (Goldman 1999, 276–271; 2001), improved by Scholz (2009). Let me first introduce a characterization of the concept of an expert. Scholz (2009, 192) condenses Goldman's basic ideas to the following definition:

“(*Expert<sub>obj</sub>-Truth*) Person A is an expert in domain D if and only if: (1) A has considerably more beliefs in true propositions and/or fewer beliefs in false propositions within D than the vast majority of people do; and (2) A possesses a set of skills and<sup>[...]</sup> methods for apt and successful deployment of this knowledge to new questions in D.”<sup>1</sup>

---

<sup>1</sup> Scholz himself thinks that this definition is too narrow, because it merely centers around the notion of truth and does not integrate the—equally important—aspects of justification, coherence and understanding (Scholz 2009, 193). I think that Scholz is completely right, but

This definition shows that there is an asymmetry between the knowledge available to the expert and to the non-expert (in what follows: the layperson). Only the expert has epistemic access to the knowledge of the domain D of expertise (Goldman calls it “esoteric knowledge”). By contrast, the layperson only has access to “exoteric knowledge”, i.e. knowledge outside the domain D (Goldman 2001, 94). This asymmetry lies at the heart of the problem of expert identifiability. The crucial question is: how can a layperson (merely with the help of the *exoteric* knowledge) identify an expert without having the relevant *esoteric* knowledge and the cognitive abilities? Goldman (2001) developed five criteria that can help the layperson to assess an expert’s reliability:

- (A) *Arguments presented by the contending experts to support their own views and criticize their rivals’ views.* Goldman distinguishes between a direct and an indirect way of evaluating the arguments of contending experts. The arguments can be evaluated directly by testing the truth of their premises and the strength of the support relation from the premises to the conclusion. The arguments can be evaluated indirectly by the *dialectical superiority*, e.g. the debating skill, the alertness and smoothness of an expert to discuss and defend her arguments against objections by others.

I think that Criterion (A) is the by far weakest of Goldman’s five criteria, because it does not accommodate the distinction between esoteric and exoteric knowledge in a satisfactory way. The direct way of evaluating the expert’s arguments is, in fact, inaccessible to the layperson, because esoteric knowledge is required to assess the truth of the premises and the strength of the argumentative support the expert can give for her arguments. The indirect way avoids this problem, but it only constitutes a rather poor indicator for expert identification. Goldman’s criterion of “dialectical superiority” primarily designates a communicative or social competence—and no *epistemic* virtue. There are, for example, professional astrologers who are quite apt at arguing for their worldview in a convincing way—which is by no means an indication of the truth of the expressed opinions or predictions. Given these considerations, I will exclude Criterion (A) from my further discussion.

- (B) *Agreement by additional putative experts on one side or the other concerning the subject in question.* Usually there is not only one expert, but many. This provides the possibility of comparing their judgments in the case of a disagreement. The most common view is that we should measure the frequency of the dissenting opinions in the relevant expert community and side with the majority.
- (C) *Appraisals of the experts’ expertise by ‘meta-experts’.* This criterion is simply a further explication of criterion (B). It makes use of appraisals to measure the credibility earned by certain experts. In fact, the scientific community has invented many operationalisations to account for the

---

I waive further explication of the concept, because the definition (*Expert<sub>obj</sub>-Truth*) offers a sufficient basis for the following considerations.

recognition of an expert in a particular field, e.g. academic degrees and awards, professional accreditations, review processes and impact factors of journals the expert has published in etc.

- (D) *Evidence of the experts' interests and biases concerning the question at issue.* There is much empirical evidence for the distortion of opinions and judgments by egoist, pecuniary or personally motivated interests and preferences. Goldman quotes a meta-analysis that compared the results of clinical trials on the efficiency of oncology drugs sponsored by pharmaceutical companies versus sponsored by non-profit organizations. Friedberg et al. (1999) found that 38% of the non-profit-sponsored trials, but only 5% of the pharmaceutical company-sponsored trials had unfavourable results.
- (E) *Using evidence of the experts' past 'track records'.* This criterion bases the possibility to evaluate the reliability of experts on the success of their predictions, advice or interventions delivered in the past. It has the advantage of being confirmed by data completely independent of the expert's knowledge. For this reason its results are epistemically accessible without any esoteric knowledge at all. For illustration, Goldman gives the example of an eclipse. An expert's prediction that there will be a total eclipse in Central Europe on 15 June 2023 is only justifiable on the basis of esoteric knowledge today; but every astronomical layperson in Central Europe can easily test its truth in ten years' time.

Before I apply the four criteria (B) to (E) to the domain of morality in order to use them to develop a precise conception of how to identify moral experts, one further comment is necessary.

At first sight, talking about 'truth' and 'justification' seems problematic in the domain of morality, because only metaethical conceptions with a realist metaphysics and a cognitivist epistemology are entitled to apply these terms to the area of ethics, i.e. in the sense of a conception of objective (realist) truth and epistemic justification. For my present purpose I do not have to interpret these concepts in this strict way. Instead I aim at integrating any metaethical accounts except radical subjectivism, i.e., all accounts that strive to capture the cognitivist intuitions that the central tenets of morality are the subject of rational moral discourse about the right and the wrong, and that reasons can be given for intersubjective agreement and disagreement. If one abandoned these conditions it would be difficult to explain why expertise could ever play a role in the area of morality. As a matter of fact most metaethical positions accept these modest cognitivist intuitions. There are some approaches trying to reconcile cognitivism with an anti-realist metaphysics (Skorupski 1999; Wright 1996). And even modern non-cognitivists—e.g. Simon Blackburn (1984; 1993; 2006) and Allan Gibbard (1990; 2003)—have developed an expressivist semantics for moral statements to “earn the right to talk of moral truth” (Blackburn 1984, 197).<sup>2</sup>

---

<sup>2</sup> For my present purposes I leave the question aside whether Blackburn's project of quasi-

In this wider sense the talk of truth and justification is in line with most modern metaethical accounts. Against this background it seems quite reasonable to apply Goldman's five criteria to the area of morality in order to investigate whether there is any suitable criterion to identify moral experts at all.

### 3. Independent Checks of Moral Expertise—Criterion (E)

I begin my discussion with Criterion (E). Goldman (2001, 106) himself stresses that, because of its justificatory independence and exoteric accessibility, this criterion is the layperson's "best source of evidence for making credibility choices". Furthermore, in the domain of empirical knowledge it has a great number of applications. The credibility of economists can be checked by the reliability of their anticipation of market trends and performance of shares, the meteorologist can be evaluated on the basis of her weather forecasts, the car mechanic by his ability to repair a car and keep it functioning and the shoemaker by his skill to make good shoes.

The problem with Criterion (E) is that, in the domain of morality, there is no epistemic resource available to conduct one of these independent checks. There are no predictions or forecasts that can be compared with actual events (McGrath 2008, 97). And there is no form of specific 'moral functioning' that could be evaluated independently of the expert knowledge of morality via exoterically accessible track records (Archard 2011, 122). On the contrary, it is possible that an economic system works very efficiently only because of moral deficiencies and that the appeal to moral norms compromises its functional qualities and effectiveness.

Many contributors to the debate have reacted to this challenge by suggesting that we do not appeal to experts of morality *tout court*, but to experts in limited or specialized domains of moral thinking (Hopkins 2007, 623–626; Jones 1999, 64–66). The leading intuition behind this suggestion seems to be that the domain of morality is far more diverse, particularized and structurally complex and that this feature of morality has to be accounted for in our theory of moral experts. Referring to 'moral expertise' *tout court* is oversimplified and fails to provide a realistic account of the role expertise actually plays in our moral discourse. In our actual practice we acknowledge moral expertise in specific areas of environmental, medical, research ethics etc., instead of expertise for morality in general. Perhaps independent checks are possible in at least some of these specified domains.

I think that it is right that we need a conceptual clarification to incorporate the phenomenon of heterogeneity in our theory of moral expertise. But as we will see, this clarification will not allow us to identify moral experts on the basis of Criterion (E).

Let me explain the relevant point in more detail. How can the particularization and structural heterogeneity of morality be explained? I think it is due

---

realism turns out to be successful—although some doubts are justified, as I have shown elsewhere (Hoffmann 2008, 188–196, 209–231).

to (a) the deep embedding of moral problems—especially of applied ethics—in empirical background knowledge, (b) the semantic complexity of so-called thick ethical concepts (like courage, envy and (medical) care), whose definitions include normative *and* descriptive conditions, and (c) the deficiencies of reasons and arguments given in colloquial moral discourses. This leads to the question: what kind of competences must an expert have to be able to rule out these complexities and deficiencies? I think they consist in the competences Peter Singer (1972) listed in his seminal paper on ‘moral experts’: the ability to argue in a precise and clear way and to detect invalid inferences, the semantic knowledge to disentangle the empirical and normative components of the meaning of thick ethical concepts—at least so far as this is possible at all—and the opportunity to “think full-time about moral issues” (Singer 1972, 117). Additionally, in many contexts a detailed and profound empirical knowledge of areas of ethical relevance is helpful, which may be one important reason for the large proportion of physicians in the medical ethics community.

I propose to ascribe ‘ethics expertise’ or ‘expertise in applied ethics’ to these experts. They are informed about the relevant developments in moral theory and philosophical ethics and have the relevant empirical background knowledge. These competences are of considerable help to give a precise account of the type of moral problems extensively discussed particularly in applied ethics. And in this respect ethics experts are superior to many laypersons, who are prone to failures in argumentation, sometimes lack empirical information about the area in question and have a sometimes insufficiently reflected grasp of the moral concepts they use. But I think it would be a misconception of their epistemic virtues to take them as the source of a privileged access to true or correct judgments about what is morally good, bad, allowed, forbidden or required. The scope of ethics expertise is limited to working out and explicating the logical structure and the empirical conditions of intricate moral problems. But it does not consist in *genuine moral expertise*, which allows the expert to have an epistemic access to esoteric *moral* knowledge.

These considerations show that ethics expertise and genuine moral expertise<sup>3</sup> must be distinguished analytically and that they differ with respect to their epistemic properties. To be precise, ethics expertise is no genuine moral expertise at all. It is characterized by far more general scientific and logical competences. So I can easily concede that the possession of these competences can—partly—be confirmed with the help of Criterion (E). This holds at least for expertise in empirical domains, e.g. concerning the effectiveness of medical treatment, environmental risk assessment etc., which can be checked on the basis of exoteric empirical knowledge via the confirmation of successful predictions. But these competences are only of indirect relevance even for the investigation of specific problems in applied ethics. For this reason I call them *secondary* competences, because they must not be identified with the epistemic competence to determine the truth or correctness of the underlying moral norms. And for the identification

---

<sup>3</sup> In what follows I use McGrath’s (2011, 126) terminology and qualify moral expertise as “*genuine* moral expertise” to avoid the confusion between the different competences of moral expertise and ethics expertise.

of experts possessing this competence, Criterion (E) is inapplicable. Hence, there are no independent checks to identify genuine moral experts.

#### 4. The Distortion of the Genuine Moral Expert's Judgment by Interests and Preferences—Criterion (D)

Criterion (D) includes the requirement that expert judgments must not be guided by self-interest or personal preferences, i.e. that they should be impartial. If we apply Criterion (D) to the domain of morality, we face a problem of justificatory circularity. This particular circularity problem occurs, because it takes moral knowledge to distinguish between the legitimate and illegitimate influence of self-interest on the genuine moral expert's judgments. Let me explain: in science, there are elaborated methodological standards to differentiate between proper validations of empirical results and illegitimate biases caused by the researcher's personal interests or preferences. In that area, *every* influence by self-interest is a potential source of bias and distortion. But the acquisition of moral knowledge is far more complicated in this respect. The crucial question is: how can legitimate and illegitimate forms of self-interest be distinguished to evaluate the reliability of a genuine moral expert's judgment?

Among ethics experts this is a highly controversial question. Richard Hare thinks that moral judgments are only 'unbiased', if they are made from an 'impersonal standpoint'. This standpoint is characterized by him as *impartial*, which means that it is preserved by any system of moral reasoning which "insists on the universalizability of moral judgments" (Hare 1981, 211). Other ethicists confirm the relevance of impartiality, but deny the close connection to universalizability (Nielsen 1985). There are also moral theories that do not contain the impartiality claim at all. The best known moral theory of this type is the theory of rational egoism, prominently discussed by Henry Sidgwick (1907, Book II) as one of the three main methods of ethics. I think that this deep disagreement among ethics experts indicates that it takes genuine moral knowledge to settle the question. For this reason, one has to refer to genuine moral expertise to solve this intricate moral problem, but Criterion (D) already assumes an impartiality claim as a presupposition to identify moral experts. Hence, the resulting justificatory structure seems to be viciously circular.

An obvious rejoinder to this problem is to refer to the distinction between exoteric and esoteric knowledge. We are no *tabula rasa* concerning morality, but we all have exoteric knowledge about morality, which is available without the appeal to expert knowledge. If the impartiality requirement is categorized as accessible via exoteric knowledge, then the justificatory circle vanishes, because laypersons have this knowledge at their disposal without any corroboration by genuine moral experts. Jones and Schroeter seem to adhere to this idea, when they speak of "wise recipients of moral testimony", who "have a sense of their own zones of comparatively greater or lesser competence" (Jones/Schroeter 2012, 228). Obviously, Jones and Schroeter qualify this meta-knowledge concerning the own moral competences as exoteric moral knowledge, too.

I see two arguments that speak against this rejoinder. First, one might suggest that it conceptualizes the scope of exoteric moral knowledge too broadly. The impartiality requirement and its moral justification are far too specialized and complicated to be part of exoteric moral knowledge. And many people who trust in moral experts are not able to give an area-specific account of their own moral competence. However, perhaps this counter-argument already admits too much. I think that the rejoinder is problematic for a second, more fundamental reason. The distinction between esoteric and exoteric knowledge it presupposes may be sharpened to a satisfactory degree in the empirical domain. But it blurs when applied to the domain of moral knowledge. That is the case, because the esoteric/exoteric-distinction has a historical aspect: as the eclipse-example illustrates, knowledge that was esoteric at a certain time can become exoteric later on. It becomes accessible to everyone by a simple observation. Therefore, the status of empirical knowledge as esoteric or exoteric may change relative to the epistemic sources available to the epistemic subject in question.

But when we look at the domain of morality and wonder where the original epistemic source of moral knowledge lies, which we view as exoteric at present, we usually have to refer to an epistemically trustworthy person, who informed us about the relevant moral rules, judgments and norms. In the early stages of our education these persons were our parents, then our educators or child care workers, and later on in school they were our teachers and tutors. Of course, not all of these people are experts in the sense explicated above (*section 2*), but they themselves acquired their moral knowledge by experts and have a status of expertise at least compared to a child's limited epistemic resources. In the end, it may turn out that *all* moral knowledge that we have acquired in the course of our life is originally based on the testimony of someone to whom we ascribed a status of expertise in questions of morality in the respective situation. For this reason, the distinction between esoteric and exoteric knowledge cannot solve the vicious circle of justification, but, at most, it obscures it by moving the decisive act of knowledge acquisition to an earlier point of time.

To sum up, these two arguments show, regardless of whether one accepts or abandons the esoteric/exoteric-distinction: Criterion (D), though helpful in the areas of empirical and even mathematical knowledge, provides no clear guidance in the domain of morality. The suspicion of a vicious circle of justification cannot be eliminated concerning this criterion to identify genuine moral experts.

## 5. The Reference to the Consensus of the Relevant Expert's Community—Criteria (B) and (C)

Two criteria are left, which both appeal to further experts to evaluate expert reliability. As mentioned above, Criterion (C) can be seen as derived from or as a specialized explication of Criterion (B). The generally approved 'meta-knowledge' to identify experts consists of measures developed by the scientific community to mutually evaluate their member's reliability. If we take a closer look at the concrete operationalisations to identify expertise, it becomes clear

that academic credentials primarily capture *the mastery of the methodological standards* of the field of expertise in question. The same holds for the peer review process carried out by established journals: it is centered around finding out whether submitted contributions are realised *lege artis*, i.e. whether and to what extent they are in line with methods evaluated as mandatory by the epistemic peers to conduct good science. Transferred to the domain of morality, these criteria primarily measure the consistency and coherence of argumentation, the knowledge of moral theory and philosophical ethics and of empirical matters that are ethically relevant. Going back to the conceptual clarification I introduced in *section 3*, Criterion (C) accounts for *ethics expertise* and not for genuine moral expertise.

This means that only Criterion (B) is left for the identification of genuine moral experts. The core idea of this criterion is to compare the judgments of different experts and to take a broad consensus (respectively the absence of expert disagreement) as an indicator for the truth or the correctness of the views in question. There is an obvious way of interpreting this idea, namely to accept a consensus theory of truth, i.e., adhere to the following idea: a moral statement MS is true by definition, if and only if the relevant expert community consents to MS's truth. But this theory of truth has to face two deep problems, which make it completely unattractive: first, the detection of moral truths becomes an empirical project of interviewing every member of the relevant expert community. Second, it is unclear who is a member of the 'relevant expert community' and who is not. This theoretical option, it seems, again leads to a circularity problem analogous to the one described in *section 3*: one has to refer to moral truth to determine the relevant expert community, but the relevant expert community's consensus defines moral truth.

So it seems reasonable to adhere to a conception of moral truth that saves the non-empirical epistemology of moral judgments and does not reduce truth to expert consensus. In this case, consensus is interpreted merely as an epistemic indicator for truth. This indicator thesis avoids the two fatal objections against the consensus theory of truth. But it only holds if two other epistemic conditions are fulfilled. In the following, I will lay out these two conditions with the help of the Bayesian approach.

The Bayesian approach in epistemology is particularly suitable for giving a precise account of some characteristics of rational belief up-dating while obtaining information by a group of epistemic witnesses or experts (Bovens/Hartmann 2003, chap. 3). I use a simple model designed by McGrew (1999) to bring out the crucial points.

Let's assume that a person has no reason for and no reason against believing a certain moral statement MS. In this situation, it is rational to apply the principle of insufficient reason (or: principle of indifference) and to assign the same degree of belief to MS as to non-MS.<sup>4</sup> Henceforth, without any further information ('a priori'), it is rational to assign the following credibilities:

---

<sup>4</sup> To calculate rational degrees of beliefs, Bayesians use probabilities. This 'calculus of epistemic credibilities' can be represented formally in analogy to the probability calculus. Hence some authors have called this representation the calculus of "epistemic probability"

$$\text{Cr}(\text{MS}) = \text{Cr}(\neg\text{MS}) = 0.5$$

It is also assumed that there are a number of experts (for the sake of simplicity: the three experts  $E_1$ ,  $E_2$ , and  $E_3$ ) who provide information about the truth of MS. We now want to know how the credibility of MS changes when the experts assent to its truth. This can be calculated with the help of a conditional credibility  $\text{Cr}(\text{MS}|\text{E})$ —in words: the credibility of MS given an expert’s positive feedback. This conditional credibility can be calculated with Bayes’s Theorem and the rule of total probability:

$$(1) \quad \text{Cr}(\text{MS}|\text{E}) = \text{Cr}(\text{MS}) \cdot \text{Cr}(\text{E}|\text{MS}) / [\text{Cr}(\text{MS}) \cdot \text{Cr}(\text{E}|\text{MS}) + \text{Cr}(\neg\text{MS}) \cdot \text{Cr}(\text{E}|\neg\text{MS})].$$

To determine the (‘a posteriori’) conditional credibility  $\text{Cr}(\text{MS}|\text{E})$  we just need to know which values we have to put into the equation for  $\text{Cr}(\text{E}|\text{MS})$  and  $\text{Cr}(\text{E}|\neg\text{MS})$ . These conditional probabilities are interpreted as measures for the expert’s reliability:  $\text{Cr}(\text{E}|\text{MS})$  expresses the probability that the expert says “MS is true!” given the truth of MS, and  $\text{Cr}(\text{E}|\neg\text{MS})$  expresses the probability that the expert says “MS is true!” given the falsity of MS. The following relation holds between these two expressions:

$$(2) \quad \text{Cr}(\text{E}|\text{MS}) = 1 - \text{Cr}(\text{E}|\neg\text{MS}).$$

Now we can calculate the *a posteriori* credibility  $\text{Cr}(\text{MS}|\text{E})$  under the presupposition of different expert reliabilities. First, I assume that all three experts are moderately reliable. I operationalise this by ascribing them a reliability of  $\text{Cr}(\text{E}|\text{MS}) = 0.6$ . Because of (2), for all witnesses the following holds: 60% indicate that an expert confirms MS if MS is true, and 40% indicate that an expert confirms MS if MS is false. With the help of rule (1), the changes in the credibility of MS based on the three expert’s assent can be calculated in three iterations.

$$(3) \quad \text{Cr}(\text{MS}|E_1) = 0.5 \cdot 0.6 / [0.5 \cdot 0.6 + 0.5 \cdot 0.4] = 0.6$$

$$(4) \quad \text{Cr}(\text{MS}|E_2) = 0.6 \cdot 0.6 / [0.6 \cdot 0.6 + 0.4 \cdot 0.4] = 0.69$$

$$(5) \quad \text{Cr}(\text{MS}|E_3) = 0.69 \cdot 0.6 / [0.69 \cdot 0.6 + 0.31 \cdot 0.4] = 0.77$$

With every further expert the *a posteriori* credibility increases. Hence, in this case the expert’s consensus can legitimately be interpreted as an indicator for truth. But does it hold generally? Let us look at the reliability of 0.5.

$$(6) \quad \text{Cr}(\text{MS}|E_1) = 0.5 \cdot 0.5 / [0.5 \cdot 0.5 + 0.5 \cdot 0.5] = 0.5$$

Regardless of the increase of the number of experts: no gain of *a posteriori* credibility is possible. This result is in line with our intuitions, because a reliability

---

(Swinburne 2001, 66–69). Personally, I prefer to stick to Carnap’s terminology and speak of credibilities (abbreviated to Cr).

of 0.5 is the reliability of a fair coin. We cannot get any additional information from it. To make the picture complete, let us finally have a look at the scenario with unreliable experts who assent more often to false than to true statements. I assume  $\text{Cr}(E|\text{MS}) = 0.4$ .

$$(7) \text{Cr}(\text{MS}|E_1) = 0.5 \cdot 0.4 / [0.5 \cdot 0.4 + 0.5 \cdot 0.6] = 0.4$$

$$(8) \text{Cr}(\text{MS}|E_2) = 0.4 \cdot 0.4 / [0.4 \cdot 0.4 + 0.6 \cdot 0.6] = 0.31$$

$$(9) \text{Cr}(\text{MS}|E_3) = 0.31 \cdot 0.4 / [0.31 \cdot 0.4 + 0.69 \cdot 0.6] = 0.23$$

Obviously, the more the number of unreliable experts increases the more the *a posteriori* credibility is diminished. The first lesson from this Bayesian model is that expert consensus is a truth indicator only if the following *reliability condition* is fulfilled: the experts have to be ‘minimally reliable’, i.e. they have to speak the truth with a probability greater than 0.5.

Furthermore, a second condition is important. The iterative calculation of the expert judgments’ contributions to the changes of the *a posteriori* credibility conducted in (3)–(5) and (7)–(9) is only applicable if these judgments are *statistically independent*. The concept of statistical independence must not be confused with the concept of independent track records discussed above in *section 3*. In that case, the track record’s independence of esoteric knowledge was decisive. But here, two expert judgments are *statistically independent* if the occurrence of  $E_1$ ’s judgment makes it neither more nor less probable that  $E_2$  judges in the same way. Usually that is the case when  $E_1$ ’s causal process of belief formation is completely unaffected by the belief forming process of  $E_2$ . An example for this kind of independence is the scenario of two eye-witnesses of an event, who have not influenced each other. Any form of interaction between the eye-witnesses (e.g. personal communication) destroys their independence. In this case,  $E_2$ ’s judgment does not contribute any additional support to  $E_1$ ’s testimony at all.<sup>5</sup> So, only if statistical independence is given, can the consensus of a big number of experts be taken as a better indicator of truth than the consensus of a smaller number.

As this Bayesian analysis shows, the claim that consensus in the relevant expert community is an indicator for truth—as calculated above—only holds if two necessary conditions are fulfilled: the reliability condition and statistical independence. The problem with these conditions is that we lack any reason to accept the reliability condition and have some reason to doubt statistical independence with respect to the domain of morality. Because all other criteria for expert identification discussed so far fail, there is no reason to ascribe minimal reliability to any individual moral expert. And without the fulfilment of this condition, the comparison with further experts exploited by Criterion (B) does not yield desirable results. Another problem is that every unreliable expert

---

<sup>5</sup> Goldman’s (2001, 98) paradigmatic example for dependence of witnesses is the guru and his slavish followers: “Whatever the guru believes is slavishly believed by his followers. They fix their opinions wholly and exclusively on the basis of their leader’s views.” For a further explication of conditional statistical independence, see Goldman 2001, 100–104.

within an expert group balances out the credibility gain of a reliable expert. If you calculate the *a posteriori* credibility for an expert group with three reliable and three unreliable experts, you will have an outcome of  $\text{Cr}(MS|E_{1-6}) = 0.5$ . The main reason to doubt the fulfilment of statistical independence was already mentioned in *section 3*: the acquisition of moral knowledge is deeply entrenched in our social practices (parental education and education in school, public discussion and controversy about moral issues, extensive interaction and scholarly exchange within expert communities etc.). Against this background, it is implausible to assign statistical independence to the judgments of individual moral experts.

So, in the end, the Criteria (C) and (B) also fall short when applied to the domain of morality.

## 6. Conclusion

My main conclusion is that all the criteria that I have considered fail to serve as a reliable information resource for the identification of genuine moral expertise. The basic problem is that these criteria are either inappropriate to identify the experts as experts concerning genuine questions of morality (Criteria (C) and (E)), that they presuppose exactly the moral knowledge we do not have available until a moral expert will inform us about it (Criterion (D)), or that necessary conditions for their application to the domain of morality are unjustified or doubtful (Criterion (B)). For this reason, the justification for the identification of genuine moral experts is either impossible or viciously circular.

I conclude by stating three further lessons from my considerations. First, we have to recognise the epistemic relevance of the distinction between genuine moral expertise and ethics expertise. The ethics expert possesses competences in science and logic that must not be equated with genuine moral competences. In the domain of morality, the ethics expert can only make justified contributions to explicate and clarify moral problems. When the moral assumptions are stated plainly and clearly, all the work the ethics expert is capable of is done. Therefore, ethics expertise is only of indirect relevance for solving moral problems. Nevertheless, the competences of the ethics expert are often helpful in addressing complicated moral questions (especially in applied ethics).

Second, the distinction also allows us to explain and to justify why moral expertise is highly established in scientific and public discourse. It can simply be explained as an appeal not to genuine moral, but to *ethics expertise*. Whether someone possesses ethics expertise, however, is testable with reference to Goldman's criteria for expert identification, because the ethics expert's knowledge and abilities are identical with a proper subset of the knowledge and abilities of scientific, medical, mathematical and engineering experts. The appeal to ethics expertise is perfectly legitimate insofar as no insurmountable problem of expert identifiability arises here.

Third, we can give a more precise account of the epistemic relevance of the 'switching' effect laid out at the beginning of my paper. The crucial point of Ingo

Schulze's observation is that we sometimes take the statement of a person not as an informative contribution to the subject matter, but as information about the attitudes and the worldview of that person. What is the epistemological reason for this switch? The Bayesian approach provides an answer: two things are important. First we have to consider the degree of reliability we ascribe to the witness in question. If we evaluate the witness as reliable (the standard assumption when ascribing a status of expertise to the witness), we will believe in the truth of her testimony. But when the trust in her reliability is undermined, her testimony does not have any credibility for us concerning the topic in question. The 'switching' effect occurs when we additionally *believe to* know the falsity of the witness' statement. In this situation, we can use the testimony to speculate about the motives for her statement—and sometimes that might contain information about her personal attitudes or worldview.

But is this interpretation epistemically justified when applied to esoteric knowledge of experts? It is only justified if the layperson does not only *believe to* know, but actually *knows* (or at least: has good reasons to believe in) the falsity of the expert's statement. These reasons are accessible to laypeople when Goldman's criteria are applicable. Hence, there may be instances, e. g. of scientific expertise, where 'switching' is epistemically justified, because the expert's reliability can be evaluated via exoteric knowledge (remember the eclipse example).

But as I have shown, in the domain of morality the esoteric/exoteric-distinction blurs and Goldman's criteria fail. We do not have any good reasons to ascribe moral expertise to anyone and—this is particularly important in the present context—we consequently do not have good reasons to dispute that someone has moral expertise. We have no epistemic means to determine the general reliability of genuine moral experts. Consequently, 'switching' is never epistemically justified in the moral domain, because justified 'switching' presupposes knowledge concerning the general reliability of an expert *and* good epistemic reasons for the falsity of the expert's judgment. However, in genuine moral matters, our only epistemic option is to discuss our disagreements point by point, question by question, argument by argument. For this reason, 'switching' in the moral domain is an illegitimate act to discredit another person's moral views, presupposing the truth of our own understanding of morality.

Let me finish with the anticipation of a possible misunderstanding of my views on moral expertise: my main thesis is purely epistemological. I do not oppose the ontological thesis that moral experts *exist*. On the contrary, I am happy to admit that there are, in fact, Aristotelian *phronimoi*. But nonetheless we have to face the result that—perhaps tragically—we are not in the epistemic position to identify them.

## Bibliography

- Archard, D. (2011), Why Moral Philosophers Are Not and Should Not Be Moral Experts, in: *Bioethics* 25(3), 119–127

- Blackburn, S. (1984), *Spreading the Word. Groundings in the Philosophy of Language*, Oxford
- (1993), Realism, Quasi, or Queasy?, in: Haldane, J./C. Wright (eds.), *Reality, Representation, and Projection*, Oxford–New York, 365–383
- (2006), Antirealist Expressivism and Quasi-realism, in: Copp, D. (ed.), *The Oxford Handbook of Ethical Theory*, Oxford–New York, 146–162
- Bovens, L./S. Hartmann (2003), *Bayesian Epistemology*, Oxford
- Friedberg, M./B. Saffran/T. J. Stinson et al. (1999), Evaluation of Conflict of Interest in Economic Analyses of New Drugs Used in Oncology, in: *Journal of the American Medical Association* 282(15), 1453–1457
- Gibbard, A. (1990), *Wise Choices, Apt Feelings. A Theory of Normative Judgement*, Oxford
- (2003), *Thinking How to Live*, Cambridge–London
- Goldman, A. I. (1991), Epistemic Paternalism: Communication Control in Law and Society, in: *The Journal of Philosophy* 88(3), 113–131
- (1999), *Knowledge in a Social World*, Oxford
- (2001), Experts: Which Ones Should You Trust?, in: *Philosophy and Phenomenological Research* 63(1), 85–110
- Hare, R. M. (1981), *Moral Thinking. Its Levels, Method, and Point*, Oxford
- Hills, A. (2009), Moral Testimony and Moral Epistemology, in: *Ethics* 120(1), 94–127
- Hoffmann, M. (2008), *Kohärenzbegriffe in der Ethik*, Berlin–New York
- Hopkins, R. (2007), What Is Wrong with Moral Testimony?, in: *Philosophy and Phenomenological Research* 74(3), 611–634
- Jones, K. (1999), Second-hand Moral Knowledge, in: *The Journal of Philosophy* 96(2), 55–78
- /F. Schroeter (2012), Moral Expertise, in: *Analyse & Kritik* 34, 217–230
- McGrath, S. (2008), Moral Disagreement and Moral Expertise, in: Shafer-Landau, R. (ed.), *Oxford Studies in Metaethics* 3, 87–107
- (2011), Skepticism about Moral Expertise as a Puzzle for Moral Realism, in: *The Journal of Philosophy* 108(3), 111–137
- McGrew, T. (1999), How Foundationalists Do Crossword Puzzles, in: *Philosophical Studies* 96(3), 333–350
- Nielsen, K. (1985), Universalizability and the Commitment to Impartiality, in: Potter, N. T./M. Timmons (eds.), *Morality and Universality*, Dordrecht, 91–101
- Scholz, O. R. (2009), Experts: What They Are and How We Recognize Them—A Discussion of Alvin Goldman's Views, in: *Grazer Philosophische Studien* 79, 187–205
- Schulze, I. (2012), *Unsere schönen neuen Kleider. Gegen die marktkonforme Demokratie—für demokratiekonforme Märkte*, Berlin–München
- Sidgwick, H. (1907), *The Methods of Ethics*, London
- Singer, P. (1972), Moral Experts, in: *Analysis* 32(4), 115–117
- Skorupski, J. (1999), Irrealist Cognitivism, in: *Ratio (New Series)* 12(4), 436–459
- Swinburne, R. (2001), *Epistemic Justification*, Oxford
- Wright, C. (1996), Truth in Ethics, in: Hooker, B. (ed.), *Truth in Ethics*, Oxford–Cambridge/MA, 1–18