



# Ethical Expertise and Moral Authority

Keith Dowding<sup>1</sup> 

Accepted: 28 March 2023 / Published online: 18 April 2023  
© The Author(s) 2023

## Abstract

Whether or not there is such a thing as moral expertise, and, if so, what constitutes it, is much debated. Empirical expertise bestows epistemic authority over propositional content; that is not the case in moral domains, technical expertise notwithstanding. This article identifies three types of agencies with some authority over decisions in moral matters. It shows that the source of the authority wielded by such agencies, while varying across the three forms identified, is based on empirical and technical knowledge and does not derive solely from any moral expertise of its members. The article examines, and finds wanting, two recent arguments for giving others authority over moral content.

**Keywords** Advice · Authority · Epistemic authority · Experts · Moral authority · Moral experts · Public policy · Testimony

## Introduction

Governments have a growing tendency to create commissions and bodies to advise on moral issues. Universities, hospitals, and private firms use internal committees or outside agencies to provide authoritative advice on medical practice and research on humans and animals, on artificial intelligence (AI) in the areas of communications, warfare, and transport. Philosophers, and moral philosophers in particular, increasingly act as expert advisors on different types of advisory committees and agencies (Schwartz 2010; Wolff 2011, 2020). Niv and Sulitzeanu-Kenan (2021) found that about a third of the philosophers they surveyed had acted at least once in such a capacity. It is increasingly suggested that society should rely upon such moral expertise: that we should endow moral experts with authority on some ethical issues as we do experts in empirical fields (Singer 1972a; Vogelstein 2015). In this paper, I will argue that while we should accept advice from agencies on moral issues, the character of the advisory evidence and nature of the authority we give to such agencies

---

✉ Keith Dowding  
keith.dowding@anu.edu.au

<sup>1</sup> Australian National University, Canberra, Australia

is not based on individuals having any superior knowledge of what is moral (what I will refer to as ‘content expertise’). I argue against two recent accounts of why we might want to give authority to others on these grounds.

The very idea of there being ethical expertise or moral experts with the authority to guide us is controversial (Archard 2011; Cowley 2005; Steinkamp et al. 2008). Many argue that expertise in empirical matters is one thing, but over moral domains another matter entirely. This view suggests that moral expertise exists not in the sense of the content of our moral views, but in providing a considered response to new problems created by emerging techniques and technologies. In fact, the major role of most philosophers in the ethical domain is to provide guidance within a regulatory framework for hospitals, universities, and other research domains. Recently, however, several philosophers have argued that moral experts can provide authoritative guidance on moral issues themselves.

In this paper, having first distinguished ‘governance’ authority from ‘content’ authority, I define three sorts of agencies that work broadly within the ethical domain. I show that they all operate with governance authority. I then consider two recent arguments suggesting we can legitimately give content authority to moral experts. Paulina Sliwa’s, I argue, shows only that on a personal level we often look to others for advice without giving them any moral authority; while David Enoch’s I find plausible only within the domain of governance authority.

## Expert Authority

We give experts such as scientists or engineers authority within their fields because of their epistemic advantage over non-experts (Raz 1990). We can trust their judgements over some proposition  $p$  outside of the content of  $p$ . In other words, they have authority because we are justified in accepting the truth of proposition  $p$ , or following the advice contained within proposition  $p$ , simply because of that epistemic advantage. Authority drawn from this form of expertise I will call ‘content authority’.<sup>1</sup> Our believing some proposition  $p$  for content-independent reasons gains credibility when recognizing that much of our knowledge derives from the testimony of others.

Generally, we accept such authoritative direction because we acknowledge that experts both have more information about the issue and are also more technically competent to assess that information. We may sometimes ask experts to explain their reasons, but often trust the specific factual information and the judgements based on their technical assessment. I will call this ‘technical authority’. For both content and technical authority, the assumption is that the propositions offered by experts are truth-apt, the content of which can be independently checked against evidence. That

<sup>1</sup> Raz calls it ‘theoretical expertise’. The authority over the content comes from the theoretical expertise of the expert. That authority is ‘content free’ in the sense that, whatever the actual content, the expert is thought to have authority over it.

is to say, in both cases we become sceptical about claimed expertise where their authoritative pronouncements do not seem to be true.<sup>2</sup>

Content-independent reasons are not the only reasons for accepting authority (Raz 1986). We can recognize that sometimes it is better to have a coordinator of collective action than no coordinator, even if we are not sure that that agent is better than we are. Raz justifies recognizing authorities as those who give us reasons to act and we do so on the grounds that following the authority is generally to our advantage. Call this ‘governance authority’.<sup>3</sup> Even in this case, however, a coordinator who does not seem to work to our advantage will lose their authority. The case against moral authority is that we never give that authority against independent evidence, only to the extent that it fits our own moral judgements.

## Morally Authoritative Agencies

Broadly speaking there are three types of agencies in which ‘moral experts’ reside: ethical commissions, ethical committees, and ethical compliance agencies. Both their work and their authoritative status vary, but all claim to offer expert ethical advice and provide moral authority.

**Ethical Commissions** (ECPs) are political or public policy bodies, often created by government, sometimes for a limited period with a specific task, but increasingly with broader, more permanent remits. One of the first was the National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research, set up in 1974 by the US Congress to make recommendations to the Secretary of the Department of Health, Education, and Welfare. It sat for four years, producing several influential reports. Other examples include, in the UK, the Animal Procedures Committee whose function—to advise the government on matters relating to animal experimentation—was made a statutory requirement by the Animals (Scientific Procedures) Act 1986; it was replaced by the Animals in Science Committee in 2013. By statute this body has 12 members, eight of whom must be doctors or vets; it must include at least one lawyer; and one member must have held a licence for animal experiments for at least six years. There is no statutory requirement for a philosopher to sit on the committee, but one usually does. A more recent UK example is the Centre for Data Ethics and Innovation. No philosopher sits on its board:<sup>4</sup> it seems philosophers are consulted less on AI issues than on medical and bioethical issues.

Philosophers figure to a greater extent on ECPs fulfilling advocacy roles—for example, the Nuffield Council on Bioethics (an independent charitable body supported by the Nuffield Foundation), the Wellcome Trust (also an independent

---

<sup>2</sup> In some fields, such as economics, it can be hard to judge, which is why some people are sceptical about the claims made in the field of economics.

<sup>3</sup> In political science ‘governance’ refers to the entire system by which society operates through the mechanisms and by which agents are held to account; it is contrasted with government which is the narrower formal institutions. Both technical and governance authority is also ‘content-free’—we give authority whatever the actual content.

<sup>4</sup> The closest being Jessica Lennard, who holds a Masters degree in Political Theory from the LSE.

charity), and the Medical Research Council (a government agency). The Nuffield Council produces reports on bioethical issues emerging from biological and medical research, and is afforded authoritative status by the UK government on bioethical matters. Three philosophers have chaired the Council since its foundation, as have three lawyers, and one former civil servant and head of an Oxford College. Many philosophers have served as committee members.

Another major area in which such commissions are active, in both public and private sectors, is AI. Important fields for AI include warfare; the Internet (in relation to both the provision and the collection of information); social and employment matters; transport, particularly driverless vehicles (though the use of AI in piloting aircraft has been important—and generally less controversial—for many years). These areas have produced fewer direct government commissions, although a host of think tanks and research institutes examines ethical issues in AI.

ECPs are typically set up by government or provide a role within government, often on a statutory basis, to advise on the creation and implementation of legislation. Such commissions are often allocated a specific task by government; at other times, as their position has become regularized, they draw issues to the attention of government; and those outside of government fill an advocacy role to persuade government that legislation is needed in certain areas of concern. Where a body is set up by government, especially when it has a statutory role, its authority proceeds from the governing and legal framework in which it operates. While expert knowledge of the medical or technical matters at issue is what leads government to respect and trust its ECP's advice, and so is in part a source of its authoritative role, that is only a small part of its authority. The link is closer to that of a principal-agent relationship, where ECPs act as agents for the government.

ECPs are tasked with a specific brief. Their judgements are backed by evidence and reasoned justifications. Their advice is usually made public—albeit often after government has considered it and adopted a position. Government uses that advice to inform its own decision processes, within both the executive and the legislature. At both stages, the executive and the legislature are subject to the pressures inherent in democracies, as other agencies, religious organizations, individuals (including other experts), and the media comment on and lobby government to follow or ignore the advice of the ECP. Rarely do governments follow the advice of ECPs lock, stock, and barrel; they sometimes take a different line altogether. The authority of ECPs is not content-independent in the sense that philosophical discussion of ethical expertise promotes—that is, that government or society accepts the content of their advice no matter what it is.

We should be wary of drawing too sharp a line between this sort of advice and other forms of expert advice to government which go through the same political and administrative processes. The authority of quasi-governmental, and especially statutory, bodies is not entirely content-based. Rather, authority is granted to them by their principals, and can be taken away; indeed, what is politically acceptable—and in a democracy that means acceptable to the media and the public—will play an important part in their considerations. Nonetheless, government utilizes the content expertise of engineers, economists, and other experts to justify their regulatory

frameworks (often well outside the purview of public discussion) in ways that do not extend to moral issues.<sup>5</sup>

ECPs are awarded a certain amount of authority, in the sense that their considerations are given weight and prominence because they are statutory bodies and respected figures serve on them. But they do not provide content-independent moral reasons for legislation or interpretation of legislation; they provide reasoned arguments. Importantly, government remains responsible for legislation. If the public disagrees with the ethics of a particular piece of legislation, the government will act (or not) on that. The difference from other sorts of expert advice is that the legislation will be judged by its content, not simply its results.

Commissions give advice, and any testimony is over facts. They provide justifications for moral arguments. Commissions generally try to come to unified conclusions, reaching their judgements through the politics of compromise, notably when given the task of making specific recommendations. This is not invariably the case; sometimes ECPs take on the role of exploring issues and providing a forum for discussion and greater understanding (Kass 2005). When they do offer recommendations, how closely government follows that advice varies, given the lobbying of other actors and political expediency. Government owns the decision and it cannot abrogate that responsibility.

**Ethical Committees** (ECAs) are bodies with a quasi-judicial role in both advising and overseeing people such as doctors, researchers, and others in areas where ethical concerns might arise, notably in human and animal welfare. Such committees operate in hospitals, universities, research institutes, government departments, and private firms. Clinical ethics consultants (CECs) also operate in medical facilities, often on a more one-to-one basis, but can be considered related, coming under the same authority relationship as ECAs. Some commercial and non-profit CECs, in fact, operate more like ethical compliance advisors.

ECAs hold much stronger authority over matters within their remit than ECPs. However, that authority does not proceed from ethical expertise. ECAs have a quasi-judicial role in both advising and overseeing what people—doctors, researchers, and others—do in the fields they supervise. Their role is to protect not only the welfare of animals and the public, but also the researchers themselves, ensuring that they do not overreach and potentially become liable to legal action. ECAs often make decisions as to whether some procedure or research can legitimately be carried out within the organization. They advise members of the organization about how they can ensure conformity with the law. They interpret and advise on the relevant existing rules and regulations. In this sense, they act like other compliance officers in organizations, such as health and safety officers, building inspectors, or committees applying rules to teaching or examination procedures in universities. While these rules concern ethical matters, the role of ECAs resembles that of other street-level bureaucrats, and their authority proceeds from the same source—that is, from their institutional, legal, and quasi-legal position.

---

<sup>5</sup> Notwithstanding the recent populist surge in demeaning experts in all fields of endeavour.

The confusion of this source of authority with that which might emanate from ethical expertise derives from the important role such committees and CECs play within hospitals. A whole branch of academia—bioethics—has arisen to consider the general issues of new technology in medical and, increasingly, in broader environmental matters. To be sure, these considerations inform ECAs and CECs as they go about their jobs, and the specifics of individual cases are why individualized advice is required. Nevertheless, these professionals operate and advise within the legal and compliance categories that govern the running of their organizations. Again, their role is that of the standard bureaucrat in any organization which is run by the rules. All rules, in any domain, require interpretation and extension as hitherto unforeseen circumstances arise.

There are many other sorts of agencies or institutions that provide recommendations to governments which involve issues of fairness and reasonable risk, such as environmental committees and food safety agencies. These can be considered as ECAs, even if not formally recognized as ethical agencies, but they also operate as many administrative bodies that make day-to-day assessments, including the interpretation of rules—and recommendations to governing authorities about changing rules. In this manner ECAs act much like many other administrative bodies, such as university exam boards or disciplinary boards or government welfare departments which assess benefit claims in terms of the rules that are set out. Such bodies are given governance authority, because we need them to make such decisions; and their decisions can, usually, be challenged by an established appeals process. They do not have to have professional ethicists on board, even though they are making ethical judgements; they operate within pre-set ethical boundaries and can usually deploy a high level of detailed specialist knowledge of technical factual aspects of the problem, often detailed in regulatory procedures. While those who study ethics in an academic setting can bring unique and important insights, they bring no further degree of moral authority.

**Ethical Compliance Advisors** (ECACs) are commercial organizations offering ethical advice to firms. Some, such as the Expert Institute and the GoodCorporation, are private companies; the Society for Human Resource Management is a non-profit organization. Such agencies generally operate alongside their clients' legal teams, to help legitimize the compliance arrangements of firms in the eyes of courts and regulatory authorities. But they also advise on turning compliance mechanisms into management tools. In other words, they advise companies on compliance with industry-wide professional norms and public expectations, thus offering input into public relations, and again often working alongside public relations firms.

The expertise of ECACs is demonstrated by their results—the success of a legal defence or the handling of a public relations crisis. They have knowledge of the professional norms of an industry. While they operate within the bounds of the ethical norms and expectations of the industry of their clients, and of the public more generally, their advice has little to do with what philosophers debate when they discuss whether or not there is such a thing as ethical expertise. It is pragmatic and legalistic advice.

Having considered three broad types of agencies which are accorded authority over propositions in the moral domain, I turn to discussion about the nature of moral

authority in the philosophical literature. I consider two important accounts: that of Paulina Sliwa and that of David Enoch. I argue Sliwa's account only accounts for personal-level advice without giving them moral authority and Enoch's works only within the domain of governance and not content authority.

## Moral Authority

Most people recognize that while philosophers have some expertise in analysing problems in the light of our morality, such expertise is not confined to those who study morality and does not give greater insight into or knowledge of moral truths (Kovács 2010; Archard 2011). While many consider that moral philosophers have an enhanced ability to identify what is morally right (Singer 1972a; Steinkamp et al. 2008; Gesang 2010; Driver 2013; Vogelstein 2015; Prialux et al. 2016), empirical studies by social psychologists suggest that moral philosophers are just as subject to biases in decision-making as other people (Schwitzgebel and Cushman 2012; Statman et al. 2020; Tobia et al. 2013a, 2013b; Buckwalter 2016).

Much of the time we make moral decisions without a great deal of thought. We simply act through morally honed dispositions. Psychologists sometimes call this 'social moral expertise', developed through our upbringing and habituated through behaviour in our society (Blasi 2005; Narvaez and Lapsley 2005, 2009). We form a moral identity adapted to the circumstances we face. That moral identity is part-and-parcel of us being moral agents (Narvaez et al. 2010). But we sometimes face moral dilemmas for which our previous experience has not prepared us. Obvious cases emerge in bioethics and AI. It is in these sorts of areas that we mostly turn to experts to help guide us through moral dilemmas. The question is what is the nature of the legitimate guidance we seek.

Here I will consider two recent arguments that suggest it is reasonable to accept the content authority of others in certain contexts. I will argue that both fail. Paulina Sliwa (2012) merely shows that we sometimes want to discuss moral issues to confirm or check our views. We do not give such advisers moral authority over the content of their advice, even if we follow it. Enoch's (2014) argument is more complex, but again its plausibility relies upon the context in which we follow such advice. It leads only to governance authority, not content authority.

Sliwa's and Enoch's arguments are based on the methods of cases—that is, the use of narratives to draw out their conclusions. Such imaginary cases or 'thought experiments' can, like real experiments, be challenged on the grounds of internal and external validity (Dowding 2022). Internal validity concerns how far we can trust the result of an experiment itself. For imaginary cases, can we conclude from internal evidence in the story what the author proclaims? External validity concerns the generalizability of findings: how far the experimental results, given their internal validity, apply more broadly outside of the experimental setting. Sliwa's narratives do not provide internal validity for her claims, but do have wider application. Enoch's also suffers from internal invalidity; and even if he does provide an existence claim, it has little external validity: it does not generalize.

Sliwa (2012, p. 176) sets up the case against deference to external moral advice by retelling Peter Singer's (1972b) imaginary drowning child case. She suggests that it would be ridiculous for me to phone my moral advisor to ask whether or not I should rescue a drowning child given I am wearing my brand-new suit. However, she then suggests that this case does not generalize to moral advice in general. She gives two examples where it seems perfectly reasonable for us to seek and rely upon others' moral advice. In one of her examples, a couple seeks advice from their siblings on the acceptability of asking the richer pair of parents to contribute more to the costs of their wedding. In another, a journalist wonders how much she should tell her family about where her editor is sending her, given the known dangers of that conflict-ridden country. Both these cases do seem to constitute real moral dilemmas, in the sense that the drowning child does not.<sup>6</sup>

However, in neither of these cases can we judge from the example alone that we can rely upon the content expertise of those from whom we seek advice. In both examples, at least part of the query to our friends and family is factual. Should the couple be honest with the parents about the differential amounts? Are the parents likely to find out, and how will they react? If the journalist were completely honest with her parents, how much anxiety would she cause them? Indeed, in these cases, as set out by Sliwa, the protagonists are looking for advice on issues where they have already, it seems, tentatively come to a decision over what they are going to do. What they seem to be doing in asking advice is seeking validation of that decision. The advice sought is deliberative. Are there other factors which they ought to have considered? While we also seek such advice from experts on other sorts of matters, there we are much more likely to expect the expert to tell us what we ought to do.<sup>7</sup>

People seek advice on moral matters for all sorts of reasons: for moral support, to stiffen their resolve to do the right thing, to be reassured that they are not doing something terribly wrong, to ensure that they are aware of all the ramifications. 'Relying on testimony' for one's moral beliefs—making up one's mind in consideration of all the facts, and discussing the ins and outs with someone one trusts—is not only morally unproblematic, it surely is the best and most common manner in which we come to difficult moral decisions. 'Taking advice' involves someone discussing a moral issue with someone else to reach their conclusion. They might be advised about social or cultural norms of which they are ignorant, or other pertinent facts—all trusted by testimony. They might have certain morally important points brought to their attention, even be taught distinctions that professional philosophers think about—what is permissible, what is one's duty, and so on. The advice leads the person to make their moral decision; they do not simply or blindly do what they are told.

<sup>6</sup> Remember that in the original Singer (1972b) version, the story is told to encourage us to see how the obvious answer in this case ought to apply to other less obvious cases, such as the plight of starving children around the world.

<sup>7</sup> In surveys people generally say that they want to make their own decisions in medical matters. However, when they are actually facing illness, people seem to want medical professionals to tell them what, if any, treatments they should have (Farrington-Douglas and Allen 2005).



At the end of the day, the betrothed couple and the journalist will have to decide what to do and live with their decision.<sup>8</sup> The point, however, is that while such advice might involve aspects of the content of the decision—factual matters about how the parents might respond, aspects of culture, or even pertinent moral considerations such as obligations and permission—these processes are not like simply trusting content expertise. They involve deliberation in which the moral agent finds their own reasons for the decision they come to in their personal dilemma.

David Enoch (2014) provides a sustained account of awarding authority or exercising moral deference to another. His argument rests entirely upon his opening example, and so I carefully examine his narrative to see whether the claims he draws from it can be sustained. I first investigate the narrative's internal validity: can he draw the existence claim for the justification of deference towards another based solely upon the content of their 'moral judgement', as we do with other forms of expertise? Enoch's target is Bernard Williams's (1995, p. 205) claim that there is no content moral expertise: a student should not simply believe a professor of ethics on some issue just because he is a professor of ethics, but needs to follow the professor's reasoning.

In order to challenge the internal validity of Enoch's imaginary case, I must quote it in some detail:<sup>9</sup>

Whenever yet another violent interaction erupts in the Middle East, Israel sees (like most other places, I think) a rallying-around-the-flag initial reaction of its public opinion. It takes some time before more critical views are voiced and allowed to be heard more clearly and loudly. In the past, I have noticed that I too seem to go with the flow on these matters. In the first few days—weeks, sometimes—I tend to see the war as in-principle justified. But often, within a couple of weeks, I come to see the war differently, indeed much more critically. And I then come to view with great shame and embarrassment my views from just a couple of weeks earlier, thinking to myself things like, 'How could I have missed how wrong this was?'. But I have this colleague—Alon, let us call him—who is different. For he—I notice—voices the moral criticism I come to endorse as the war proceeds much earlier, indeed, from its very outbreak. And now ... another armed conflict erupts. I think about it, familiarize myself as much as possible with the relevant facts, and it seems to me that we are by and large in the right. But I talk to Alon, who once again tells me how wrong this whole affair is. Let us further suppose that I do not have to worry just about forming my opinions. Rather, I have to *do* something, something that will have an effect on the real world ... say I have to cast a deciding vote on funding for the war ... Here it seems to me, there is considerable pressure

<sup>8</sup> Alison Hills (2009) distinguishes moral testimony from moral advice. What we call such advice is not really important; the issue is more that in these examples, detailed factual evidence or views on facts are an important part of the advice.

<sup>9</sup> Given the personalized nature of the example, it might be something Enoch has noticed about himself, but that merely invites closer empirical analysis of the manner in which Enoch's own thought processes work and the factors that lead to his opinions. I will assume it is entirely imaginary.

to say that given that Alon—as I know, or anyway wholeheartedly believe—was much more often right than I was about such matters in the past, I should believe that the war cannot be morally defended, and act accordingly. (Enoch 2014, pp. 230–231)

Now, we can specify in advance that if ‘I’ do have to vote on the war, given this example, that ‘I’ am morally required to vote against the war. Furthermore, we can agree with Enoch that this judgement is based on the need not to place people at a ‘higher-than-need risk of being wronged’ (Enoch 2014, p. 232). In other words, we fully accept that ‘I’ wholeheartedly agree that Alon is more likely to be right in his beliefs about the moral case against the war than my views at the time of the vote. Call this time  $t_1$ . The issue at stake is not whether I should (am ‘morally required to’, as some would have it) defer to Alon at time  $t_1$ —we agree on that—but rather is the nature of the grounds on which I ought to follow his beliefs rather than mine at time  $t_1$ .

There are several grounds on which I might defer to Alon: (1) he knows more about the facts of the war than I do; (2) he is able to process whatever facts we both have about the war better than me; (3) he is not swayed by irrelevant contextual matters—that is, he does not get swept up in the emotions as I do (and this might help explain point 2); (4) I simply note that my final views on such wars (at time  $t_2$ ) always agree with his at time  $t_1$ .<sup>10</sup>

None of (1) to (4) constitutes the type of moral deference that Enoch requires and he argues against each being the reason why ‘I’ should defer to Alon. Enoch requires that Alon is a superior judge of the ‘content’ of the moral decision than ‘I’ am. He says, ‘the expert is much more likely to be right on the relevant question’ (Enoch 2014, p. 233). To discount (1) as a reason to defer to Alon, Enoch stipulates that ‘there are no non-moral facts that Alon knows of that I do not’ (Enoch 2014, p. 233). In response to the reply, ‘then on what grounds should we trust Alon?’, Enoch suggests that ‘I’ can wholeheartedly believe Alon is more likely to be right because ‘evidence of past record is relevant’. This evidence can be as in the extended example; or that ‘I’ have *other* evidence that Alon is more likely to be right; or that Alon is more reliable in other domains where ‘I have reason to think morality is not different from them in this respect’ (Enoch 2014, p. 233).

By the narrative, the evidence that Alon is right at  $t_1$  is contained in considered judgements I come to. Enoch suggests that at  $t_1$  I am in a state of moral uncertainty. This moral uncertainty exists because although at  $t_1$  I think the war is justified (I do not seem to be uncertain), I am aware that Alon disagrees, and that I always come to agree with him. Thus, it seems, by the narrative Enoch offers, the only reason we can plausibly have for voting against the war at time  $t_1$ , despite thinking it justified, is that Alon says the war is unjustified and Alon has proved in the past to be a better judge of my final considered moral judgement. Thus, in voting against the war I am (reasonably enough) betting that Alon is right, and that I will come to see that

<sup>10</sup> I assume, given Enoch’s account, that what I come to decide at  $t_2$  is my final considered moral judgement and I do not change my mind again.

in time. If, on this occasion, I do not come round to Alon's views, then I will regret my decision. I am thus deferring only to Alon's ability to predict at time  $t_1$  what my considered moral judgement will be (which I will not discover until I have overcome the 'rally-round-the flag' emotion) at time  $t_2$ . I am not deferring to his superior moral knowledge, but to his psychological propensity, at a given time, to overcome strong emotion. I do not think he has content authority, since the judge of his ability to judge is also mine.

What other evidence is possible in such cases (where, remember, evidence is not based on any factual information Alon has)? What grounds could one have to see that Alon is more reliable than I am in other domains where morality is not different? Without examples, one can only see those contained in (2) and (3) above. Alon is able to process information better, especially when subject to irrelevant contextual factors such as peer pressure and emotions. The plausibility of Enoch's extended example relies on (2) and (3). It relies upon 'I' being swayed early in the war by nationalistic emotional responses, which over time 'I' overcome. By time  $t_2$ , I am able to process the information that Alon managed to do at time  $t_1$ .

One question that Enoch ponders is the nature of Alon's superior conclusion. Do we require there to be a 'realist' or 'objectivist' account of moral evaluations? Enoch suggests he does not need such an account; and he does not. At least not in his extended example, where the justification of Alon's superiority is that the considered judgement of 'I' corresponds to Alon's. The 'other' information in the above paragraph could be the truth of the objective moral evaluation demonstrable by some evidence outside of the judgement of Alon and 'I'. Such a claim would assume that moral expertise is like other expertise, and we can simply trust conclusions without understanding them because we can discover in *other ways* that they are correct.

It seems, then, by the thought experiment Enoch offers, that the only reason we can plausibly have for voting against the war at time  $t_1$  is that Alon says the war is unjustified and Alon has proved in the past to be a better judge of the considered moral views of 'I' at time  $t_1$  than 'I' am. The only plausible grounds, given that Alon does not have any factual information at  $t_1$  that 'I' do not have, and given there is no independent check against objective moral reality, is that Alon is able to process the information we both have in a manner not subject to irrelevant emotions. There is nothing controversial in that claim. Many people can realize that they should not make ethical decisions in heightened emotional situations—moments of anger, jealousy, or indeed jingoism. However, that does not show Alon has content expertise; it just shows that he is, at time  $t_1$ , better technically equipped than 'I' am to process moral judgements. Bernard Williams would not disagree. His argument is against content moral expertise based on an inability to understand the reasoning of the professor, not the inability, when subject to strong emotions, to utilize such reasoning.

So, all we can conclude from the argument of Enoch is that I need to be aware that, at time  $t_1$  Alon's judgement is different to mine at  $t_1$ , that I follow his judgement rather than mine, and that at some future time,  $t_2$  I come to see his judgement at  $t_1$  was better than mine at  $t_1$ . That requires that I come to my own judgement at  $t_2$  that is independent of my accepting his judgement at  $t_1$ . In other words, in order for me to justifiably defer to Alon's moral judgements at any time ( $t_1$ ) I must still be capable of coming to my own judgements at some time ( $t_2$ ) to make the judgement that

his judgements at the earlier times ( $t_1$ ) are better than mine at that time. This is so because moral expertise is based on justificatory ('technical') evidence not simply conclusion tested against empirical evidence.

In the ethical case, all we have is the technical expertise as it is displayed in the reasoned assessments of the experts. Thus, it cannot be by any assessment of the content of the judgements themselves that we come to value the relative expertise of the philosophers, but rather by how we judge the technical expertise they display. In that sense, we cannot give them content-independent authority on the basis of their reasoned views; rather, we judge them by the content of how they provide their reasoned opinions. To the extent that we follow the conclusions of moral experts, we do so on the grounds that we find the content more acceptable, not on content-independent grounds (though we might admire the way that a position is developed, even if we reject its conclusion).

## Conclusion

An extensive debate addresses whether or not there is such a thing as moral expertise and, if so, in what that expertise consists. Whether and in what sense philosophers have moral expertise matters, given their increasing presence in various public policy domains. I first considered how expertise can be thought to give authority, in the sense of providing us with content-independent reasons for believing or doing something. We might follow expert advice simply because we believe that the expert is more likely to know the correct or best course of action. We do not ordinarily believe that simply on the basis of someone being correct—we might think someone making predictions has been right in past through luck—but because we think they have some technical skills that lead them to those better conclusions. In this sense expertise does not simply consist in being more likely to know the correct or best course of action, but in the knowledge and, particularly, skills that experts develop that lead them to be correct more often than non-experts. Such expertise can exist in areas that are not, or not straightforwardly, truth-apt.

Enoch's case is one where, previously, the truth-aptness of Alon's judgement at  $t_1$  has been tested by the judgement of 'I' at later times  $t_2$ , and so is trusted at this time  $t_1$ . The internal validity of this particular thought-experiment is provided by the explanation that Alon is not subject to nationalistic emotions in the manner of 'I' or some other, less specified reasons ('trust in other areas'). If the example has any external validity, it will be on the same grounds. This is not content moral expertise, but rather a trusted psychological disposition—a 'cool head'. Similarly, we might trust agencies to provide less emotional advice on moral matters in the public or political domain, without thereby awarding them content expertise. Rather, we come to judge that their advice is dependable.

The Alon case can be contrasted to the narratives offered by Sliwa, where the nature of the consultation is more deliberative. We often seek advice on moral matters, but we do so to ensure that we have not overlooked any moral issues or factual matters, or simply to be reassured we are doing the right thing or, at least, not doing something dreadfully wrong. In the personal domain, we do not hand over

content moral authority. We do not do so because we view moral agency in terms of responsibility. We expect moral agents to take responsibility for their actions, and not simply to defer to the opinions of experts. As Bernard Williams argues, we expect them to be able to produce reasons for their moral decisions. In the Alon case, a moral reason is provided if one were to say, 'I followed Alon's advice; I do not trust my judgement given the emotive circumstances, and his views have proved sage in the past'. But that is just to say one does not trust one's moral agency in these circumstances.

When we turn away from abstract considerations about ethical expertise to what moral experts actually do, we see that their recommendations are often not the major aspect of their role. For example, clinical ethics consultants explain to patients the options offered by their health-care providers, or they may help resolve conflict between family members and health-care providers over how to treat patients, particularly when those patients are not in a position to provide informed consent (Fiester 2012). Indeed, these skills are codified by professional associations such as the American Society for Bioethics and the Humanities (ASBH 2011; Kodish et al. 2013; see Iltis and Sheehan 2016, p. 429).

Government commissions make recommendations about how the law should be framed in areas of medical practice or scientific research where new techniques open up possibilities not previously foreseen. Or they sit on committees within firms advising about issues to do with AI, robotics, or social media. Such recommendations from committees do not simply tell governments or managers what to do, but are accompanied by a set of justificatory reasons or arguments. These justificatory reasons are often compromises between conflicting views within the committee and take practical and technical considerations into account (Wolff 2011). The work of such committees is largely factual, concerning processes that new technology opens up. It is the detail and the context that are important in these recommendations (Wolff 2020). Ethical expertise in this sense is command of the moral underpinning of the recommendations given new techniques, and often knowledge of the consensus within the discipline about the welfare possibilities opened up by new techniques, as well as some knowledge of public opinion on such matters (Rasmussen 2016). Thus, the knowledge of moral experts is often contained in the knowledge they gain from examining the specific evidence for the complex cases about which they are asked to provide advice, as much as the technical competencies that Singer (1972a) and others point to. That is not to downplay those technical competencies; but they are directed at justifying the conclusions that these agencies reach for consideration by those whom they advise.<sup>11</sup>

I distinguished three types of moral agencies. They are rather different, but in diverse ways we award these agencies governance moral authority. ECPs provide specific detailed advice, often technical and factual, to government, usually made public for comment and action within normal democratic political processes. Their authority is that devolved to agents within a legal advisory framework; they have no

---

<sup>11</sup> O'Neill (1996) offers a nice short account of how the expertise of moral philosophers can be used in bioethics.

decision-making powers, but rather provide relevant information and justification for the course of action they advise. Some play advocacy roles, but do so also by providing technical and factual information.

ECAs operate as normal bureaucrats, interpreting and implementing the rules of the organization of which they are a member. Their work involves decisions which often have important ethical dimensions; again, it is their relevant experience and knowledge of previous precedents and the legal framework which provide their authority. Outside of regulatory procedures, they may provide advice, to researchers or doctors for example, but the final responsibility for decisions rests with those researchers and doctors.

ECACs are commercial agencies that provide quasi-legal, public relations, and compliance advice to companies within the sphere of moral issues. They have no special authority other than their experience in such matters.

The three types of agencies I have identified all play a role in providing expert advice on ethical issues. Their roles and the type of advice they offer differ, however, as does their authoritative status. The issues related to the nature of ethical expertise debated in much of the academic literature is largely separate from the issues concerning the authority of agencies in the real world that provide services within the domain of ethics.

**Acknowledgements** This paper was first given as an (online) public lecture in Oslo, and I thank the audience for comments. I would like to thank Anne Gelling, Eilev Hegsted, Stephen King, Lars Moen, Yarden Niv, and Alex Oprea for their written comments, and especially Jakob Elster and Cathrine Holst for their extensive comments and suggestions.

**Funding** Open Access funding enabled and organized by CAUL and its member institutions.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Archard, David. 2011. Why moral philosophers are not and should not be moral experts. *Bioethics* 25: 119–127.
- ASBH. 2011. *Core competencies for healthcare ethics consultation*. Glenview, IL: American Society for Bioethics and Humanities.
- Blasi, Augusto. 2005. Moral character: A psychological approach. In *Character psychology and character education*, ed. Daniel K. Lapsley and F. C. Power, 67–100. Notre Dame: University of Notre Dame Press.
- Buckwalter, Wesley. 2016. Intuition fail: Philosophical activity and the limits of expertise. *Philosophy and Phenomenological Research* 92 (2): 378–410.
- Cowley, Christopher. 2005. A new rejection of moral expertise. *Medicine, Health Care and Philosophy* 8 (3): 273–279. <https://doi.org/10.1007/s11019-005-1588-x>.

- Dowding, Keith. 2022. How to use imaginary cases in normative theory. *Metaphilosophy* 53 (4): 512–525.
- Driver, Julia. 2013. Moral expertise: Judgment, practice, and analysis. *Social Philosophy and Policy* 30 (1–2): 280–96.
- Enoch, David. 2014. A defence of moral deference. *Journal of Philosophy* 111 (5): 229–258.
- Farrington-Douglas, Joe, and Jessica Allen. 2005. *Equitable choice for health*. London: IPPR.
- Fiester, A. 2012. The ‘difficult’ patient reconceived: An expanded moral mandate for clinical ethics. *American Journal of Bioethics* 12 (1): 2–7.
- Gesang, Bernward. 2010. Are moral philosophers moral experts? *Bioethics* 24 (4): 153–159.
- Hills, Alison. 2009. Moral testimony and moral epistemology. *Ethics* 120 (1): 94–127.
- Iltis, A., and M. Sheehan. 2016. Expertise, ethics expertise, and clinical ethics consultation: Achieving terminological clarity. *Journal of Medicine and Philosophy* 41 (2): 416–433.
- Kass, Leon R. 2005. Reflections on public bioethics: A view from the trenches. *Kennedy Institute of Ethics Journal* 15 (3): 221–250.
- Kodish, E.J., et al. 2013. Quality attestation for clinical ethics consultants: A two-step model from the American Society for Bioethics and Humanities. *Hastings Center Report* 43 (1): 26–36.
- Kovács, József. 2010. The transformation of (bio)ethics expertise in a world of ethical pluralism. *Journal of Medical Ethics* 36 (12): 767–770.
- Narvaez, Darcia, and Daniel K. Lapsley. 2005. The psychological foundations of everyday morality and moral expertise. In *Character psychology and character education*, ed. Daniel K. Lapsley and F. C. Power, 140–165. Notre Dame: University of Notre Dame Press.
- Narvaez, Darcia, and Daniel K. Lapsley. 2009. Moral identity, moral functioning, and the development of moral character. In *The psychology of learning and motivation*, ed. Daniel M. Bartels, Christopher W. Bauman, Linda J. Skitka, and Douglas L. Medin, 237–274. Burlington, VT: Academic Press.
- Narvaez, Darcia, T. Gleason, and C. Mitchell. 2010. Moral virtue and practical wisdom: Theme comprehension in children, youth, and adults. *Journal of Genetic Psychology: Research and Theory on Human Development* 17 (4): 363–368.
- Niv, Yarden and Raanan Sulitzeanu-Kenan. 2021. An empirical perspective on moral expertise: Evidence from a global study of philosophers. *SSRN ELibrary*. [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=3922183](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3922183).
- O’Neill, Onora. 1996. Medical and scientific uses of human tissue. *Journal of Medical Ethics* 22 (1): 5–7.
- Priaultx, Nicky, Martin Weinel, and Anthony Wrigley. 2016. Rethinking moral expertise. *Health Care Analysis* 24 (4): 393–406.
- Rasmussen, Lisa M. 2016. Clinical ethics consultants are not ‘ethics’ experts—but they do have expertise. *Journal of Medicine and Philosophy* 41 (3): 384–400.
- Raz, Joseph. 1986. *The morality of freedom*. Oxford: Oxford University Press.
- Raz, Joseph. 1990. Introduction. In *Authority*, ed. Joseph Raz, 1–19. Oxford: Basil Blackwell.
- Schwartz, Jason L. 2010. Lessons and opportunities for presidential bioethics commissions. *The Good Society* 19 (1): 10–15. <https://doi.org/10.1353/gso.0.0096>.
- Schwitzgebel, Eric, and Fiery Cushman. 2012. Expertise in moral reasoning? Order effects on moral judgment in professional philosophers and non-philosophers. *Mind and Language* 27: 135–153.
- Singer, Peter. 1972a. Moral experts. *Analysis* 32: 115–117.
- Singer, Peter. 1972b. Famine, affluence, and morality. *Philosophy and Public Affairs* 1: 229–243.
- Sliwa, Paulina. 2012. In defense of moral testimony. *Philosophical Studies* 158: 175–195.
- Statman, Daniel, Raanan Sulitzeanu-Kenan, Micha Mandel, Michael Skerker, and Steven De Wijze. 2020. Unreliable protection: An experimental study of experts’ in-bello proportionality decisions. *European Journal of International Law* 31 (2): 429–453.
- Steinkamp, Norbert L., Bert Gordijn, and Henk A. M. J. ten Have. 2008. Debating ethical expertise. *Kennedy Institute of Ethics Journal* 18 (2): 173–192.
- Tobia, K., W. Buckwalter, and Stephen Stich. 2013a. Moral intuitions: Are philosophers experts? *Philosophical Psychology* 26: 629–638.
- Tobia, K., G. Chapman, and Stephen Stich. 2013b. Cleanliness is next to morality, even for philosophers. *Journal of Consciousness Studies* 20: 195–204.
- Vogelstein, Eric. 2015. The nature and value of bioethics expertise. *Bioethics* 29 (5): 324–333.
- Williams, Bernard. 1995. *Making sense of humanity*. Cambridge: Cambridge University Press.
- Wolff, Jonathan. 2011. *Ethics and public policy: A philosophical inquiry*. London: Routledge.

---

Wolff, Jonathan. 2020. Method on philosophy and public policy: Applied philosophy versus engaged philosophy. In *The Routledge handbook of ethics and public policy*, ed. Annabelle Lever and Andrei Poama. Abingdon: Routledge.

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.