Synthese (2025) 206:294
https://doi.org/10.1007/s11229-025-05382-z

ORIGINAL RESEARCH

How moral philosophers can help society

Kian Mintz-Woo[1,2]

Received: 21 February 2025 / Accepted: 13 November 2025

Abstract

This paper argues that moral philosophers can have a special role in helping members of society come to choose which moral theories to believe. Importantly, the argument does not depend on the idea that moral philosophers (more) reliably have true moral beliefs (or are "Strong Moral Experts"). Instead, the argument is that moral philosophers are well-placed to develop understanding of moral theories by drawing out valid implications (they are "Weak Moral Experts"). By developing valid moral arguments, and by making the relevant implications accessible to society, moral philosophers can help people understand the costs and benefits of various moral theories, allowing them to make more informed choices. This does not imply that everyone will agree; there is room for disagreement about the weight to put on various theoretical costs and benefits. But it does give a metaphilosophical picture of the role of moral philosophers, justify certain kinds of public philosophy, and explain the value that moral philosophers can add to society at the philosophy-public interface.

Keywords Ethics of belief · Metaphilosophy · Moral expertise · Philosophers · Philosophical methodology · Public philosophy

✉ Kian Mintz-Woo
mintzwoo@iiasa.ac.at

[1] International Institute for Applied Systems Analysis, Laxenburg, Austria

[2] Department of Philosophy | Sustainability Institute, University College Cork, 4 Elderwood, Cork, Ireland

Published online: 08 December 2025

1 Introduction

What should the relationship between philosophers and the non-philosopher public be in the formation of moral beliefs?[1] Roughly speaking, there are three families of answers.[2]

Perhaps the first family of answers is that there should basically be no relationship; philosophy can be done more or less without significant or at least intentional interaction with the public at all. Philosophers need not consult members of the public; the public has no special need to know what philosophers are doing. Perhaps philosophers should attend to the way members of the public use *words* instead of stipulating or claiming unrelated philosophical meanings for those words.[3] But with this family of answers there is no especially important relationship between philosophers and the public with respect to the formation of moral beliefs.

This family of answers has recently and notably been challenged by experimental philosophers, whose goal is to empirically determine what the philosophical beliefs of the public are (a prominent example is Knobe, 2003). This is often in service of suggesting that philosophers' reliance on some number of intuitions is either suspect or at least not reflective of the way the public forms beliefs (Knobe, 2007; Horvath & Koch, 2020).[4]

A second family of answers is that philosophers can benefit from the public, such that interaction with the public improves the formation of moral beliefs. A recent proponent of this type of position is de Shalit (2020), who argues that political beliefs and theories amongst philosophers should incorporate insights or ideas from members of the public. He argues that philosophers should intentionally engage members of the public to elicit insights or ideas which can be considered alongside philosophical positions as input into a "public reflective equilibrium" process (cf. §3). The idea is that discussing political and social issues with members of the public can help provide material for theorizing (also cf. Archard, 2011; Wolff, 2020).

[1] This presupposes a distinction between academic or professional philosophers and a non-academic public. This distinction is relatively recent in the history of philosophy, so this formulation of the question does not apply when this distinction was less applicable. I have in mind what McIntyre (2022) suggests, whereby Kant helped usher in professional philosophy as something done by academics and not laypeople (also cf. Weinstein, 2014). If we accept that historical account, this question would only arise in this form in a contemporary, roughly post-Kantian, period.

[2] I believe this is an especially challenging version of the general question of the role of (academic) expertise in society. For instance, it is considerably less controversial that members of the public should broadly defer to scientific experts in society, e.g., take the most recent reports of the Intergovernmental Panel on Climate Change to reflect the most justified beliefs about climate change. Areas of social science may be more complex; perhaps the aims of society are the role of the public and the achievement of those aims should be guided by academic expert knowledge (Christiano, 2012). But for issues of values and morality, it is considerably less obvious what the relationship between philosophers and the public should or could be. For similar reasons, I am interested in how individuals form beliefs—not, for example, in social questions like how the public should select laws or policies.

[3] Moves of this sort are made by ordinary language philosophers (a recent explanation of the resultant methodology is made by Baz, 2016).

[4] Doris (2021, p. 68–69) puts this nicely: "systematic empirical works allows access to responses other than that of the philosopher, or small group of philosophers, with a theoretical investment in the case… the beliefs of ordinary folk tell us more about folk beliefs than the beliefs of professional philosophers."

Of course, this kind of method may fail to be representative or systematic (Mintz-Woo, 2020), but it is clearly an attempt to improve philosophers' moral belief-formation by intentionally interacting with the public.

However, while not staking a position on whether either of these families of answers will succeed, I want to develop an account in the third family of answers, since I believe this family of answers is less commonly defended. In this family of answers, the public could benefit from philosophers when forming moral beliefs.

One obvious way that this family of answers could be endorsed would rely on the claim that moral philosophers (relatively reliably) form true moral beliefs so the public would benefit if the moral philosophers told them which beliefs are true. However, there are important challenges to this kind of position. One is that, despite having access to the same kind of information, there is widespread disagreement in moral philosophy (Bourget & Chalmers, 2014, 2023); some think that this licenses skepticism that we should endorse moral claims when similarly situated interlocutors disagree (for a discussion of this debate, cf. Cosker-Rowland, 2017). Another challenge relies on the claim that one does not get full moral credit if one believes or acts on the (mere) basis of moral testimony (Hills, 2009). Furthermore, some philosophers argue that beliefs formed (merely) on the basis of moral testimony (perhaps unlike other forms of testimony) never qualify as knowledge (Hills, 2013). However, methodologically, the question is how far someone defending the view that philosophers can contribute to society can get while assuming *as little as possible*. Obviously, if these various challenges can be met, a more ambitious claim can be defended—but the purpose of this project is to see what can be done under the most pessimistic assumptions.

Is there a way of rescuing the position that philosophers have something important to add to the moral belief-forming practices of the public even if we grant (at least arguendo) that simply telling the public which substantive moral claims to believe is either inappropriate or hubristic? This paper argues that the answer is "yes".

The key idea is that moral philosophers are well-placed to explain which inferences or conclusions validly *follow* from substantive normative (and non-normative) premises. Equivalently, we could say that moral philosophers can convey valid *arguments* (understood in a very wide sense) which allow a member of society to better *understand* various normative positions (for simplicity, I call combinations of normative theories, principles, and judgments "(normative) positions", with a primary interest in moral theories).[5] The point is not that philosophers would just explain their theories, since the public usually has only "a very partial appreciation of the theories" (Christiano, 2012, 46). Nor is it that philosophers would just give their considered practical judgments. Instead, they would explain what can validly be inferred from normative positions, inferential links that may be subtle or non-obvious to those not steeped in the philosophical nuances. This allows members of the public to compare alternative normative positions, together with their sets of implications, in a more

[5] Throughout, I use the term "valid" inclusively. Traditionally, it refers to structured arguments where the conclusion follows necessarily from the premises. However, while these narrowly *deductively* valid inferences are an ideal form for philosophical reasoning, they do not exhaust the genre. I intend abductively and inductively valid inferences to be part of my valid set, at least for contexts where different forms of validity are appropriate (Peirce, 1931–1958).

nuanced and comprehensive manner, than they could in the absence of this philosophical input. This also suggests an interesting metaphilosophical picture: a potential role for (moral) philosophers is to allow members of society to better *understand* the moral commitments that various positions imply (understanding as the point of philosophy is explored carefully by Hannon & Nguyen, 2022).

The argument proceeds in two stages. The first stage is to establish that philosophers are well-equipped to provide this kind of information. In order to do that, it is necessary to carefully explain what philosophers can add—and what philosophers are especially *well-placed* to add—to the (moral) belief-formation practices of members of the public. In other words, we need an appropriate account of *moral expertise*. In §2, I introduce two forms of moral expertise, precisifying and clarifying a distinction that is important in the moral expertise debate. The novel contribution is distinguishing the two forms of moral expertise (weak or strong moral expertise) by which type of moral arguments they are meant to produce (valid or sound moral arguments, respectively), together with providing some reasons for thinking that moral philosophers are generally well-placed to act as weak moral experts. The second stage is to indicate that this kind of information is, or could be, valuable to members of the public. In §3, adopting the role of weak moral expert for moral philosophers, I show that this account of moral expertise suggests an interesting metaphilosophical or public philosophical contribution that philosophers can make to the public. They can explain what valid inferences follow from normative positions, thus increasing understanding of these positions. I discuss a couple objections in §4. Finally, in §5, I conclude.

2 Which form of moral expertise should we take philosophers to have?

In this section, I consider three positions to take in the discussion of philosophical expertise with respect to moral claims: first, that there is no relevant philosophical expertise; second, that there is Strong Moral Expertise; and third, that there is Weak Moral Expertise. The third position avoids some issues associated with the first two, but also provides some reasons for believing that moral philosophers are especially well-placed to play the role of (Weak) Moral Experts, compared to non-philosophers.

One position in the discussion of philosophical expertise is a skeptical one. Nielsen, for instance, denies that there is philosophical expertise, since the ability to think carefully is not limited by profession:

> there is nothing in the way of philosophical expertise, technique, knowledge or bright new analytical tools which they trot out or which are available to the philosopher which will enable her to get a purchase on these [moral or social justice] problems in the way that a political scientist, a literary critic, a novelist, a lawyer, a sociologist, a political economist or a historian cannot or can only with a kind of difficulty from which the properly equipped philosopher is free. (Nielsen, 1987, p. 14)

Nielsen is in an important sense right; any skill with understanding or applying normative positions need not exclusively be the realm of philosophers (also cf. Wolff, 2020). However, it is less controversial that moral philosophers are generally trained to know more than other colleagues—whether inside or outside of the academy—*about* moral theories, including how to wield them (e.g. Crosthwaite, 1995).

To understand the import of this point, it is helpful to distinguish two forms of moral expertise, drawing on the applied (especially biomedical) ethics literature on moral expertise: "Weak" and "Strong Moral Expertise".[6] While the various definitions in this literature are not precisely co-extensive with mine in all cases, I believe the definitions I provide below help both to simplify the extant discussions and to unify key elements from the literature in a systematic way (Birnbacher, 2012, Nickel, 1988, Yoder, 1998, Zoloth-Dorfman and Rubin, 1997).

Let us set aside a highly contingent and, I would suggest, irrelevant mark of philosophical expertise. It is not essential to have structured "time to think and reflect about ethical issues" (pace Singer, 1972, 1988, p. 154). If two people come to the same conclusions from the same ethical starting points, and one required more time to do so, that's not a mark of greater expertise.[7]

A helpful way of thinking about the distinction I will draw between two types of moral expertise relates to logical argumentation. In terms of argumentation, Strong Moral Experts know the morally and non-morally correct *premises*, and can soundly deduce true moral *conclusions*. In contrast, Weak Moral Experts know of some *potentially* true moral and non-moral premises, and can validly deduce *potentially* true moral conclusions (cf. Birnbacher, 2012).[8] This parallels the familiar distinction between sound and valid reasoning where soundness requires the truth of the premises and necessary truth preservation between the premises and the conclusions whereas validity requires only the necessary truth preservation.[9]

[6] Confusingly, these forms are often called "ethical/moral" or "semi/full" or even "narrow/broad" expertise (e.g. Steinkamp et al., 2008; Gesang, 2010; Priaulx et al., 2016; Hegstad, 2023). All of these terms pick out similar (or even the same) distinctions, so it is unfortunate that we do not have completely consistent nomenclature. I believe a terminological distinction between "Strong Moral" and "Weak Moral" expertise more perspicuous, so I adopt it here. To add to the confusion, some authors even use terms like *ethics expertise* as a *family* term to include both "moral" *and* "ethical" kinds of expertise (Rasmussen, 2005).

[7] Indeed, it seems to me that, insofar as a given potential expert is good at reliably drawing the valid inferences from some normative position, it would be an indication of *greater* moral expertise if she were capable of doing so quickly and so did not require time to think and reflect. Analogies from other fields of expertise help: an expert cook can work more quickly and effectively than a new cook; an expert chess player can analyze a position with greater speed and insight than a novice. Here, it appears that Singer has confused a commonly helpful precondition for *developing* the relevant forms of knowledge with a condition for *having* them.

[8] Throughout, "potentially correct" is elliptical for "taken to be potentially correct", which can be contextually specified as relative to a time, place, or group as appropriate. For our purposes, I am primarily thinking of which theories are taken to be potentially correct by the contemporary philosophical community (e.g. for moral philosophers, perhaps deontology, consequentialism and virtue ethics, and derivatives thereof), especially—but not necessarily only—the professional community as expressed through published positions. This is largely on pragmatic grounds; no assumptions about truth-tracking are to be inferred.

[9] Here again, I am taking deductive reasoning to be the illustrative case (see fn 5); parallel claims can be made for inductive or abductive argumentation (e.g., "cogent" for "sound", etc.).

This provides a conception of *Strong Moral Experts* as follows:

Definition 1 *(Strong Moral Experts). Those who know (or at least reliably believe) the correct normative position. Furthermore, they are able to (at least reliably) validly infer the correct and important morally relevant conclusions from the correct normative position and other relevant non-normative facts.*

This tripartite structure combines a normative epistemic component and a non-normative epistemic component with a skilled component. In short, a Strong Moral Expert (reliably) infers the morally relevant conclusions and does so non-accidentally, inferring from the true normative position using relevant non-normative facts. She does not waste time drawing trivial or irrelevant conclusions, but tends to infer conclusions that are important in practice.

However, Strong Moral Expertise is too strong in this context for philosophical expertise on (at least) two grounds. First, for Strong Moral Experts, there is no reason to reject Nielsen's skepticism: a political scientist or a literary critic might adopt the correct normative position just as reliably as a moral philosopher. More importantly, even if one group systematically did develop the right beliefs more often, there is no generally reliable way to ascertain which group did (e.g. Cholbi, 2007).[10] Second, if the assumption is granted that moral testimony is objectionable, then simply reliably adopting the true normative position is not helpful for a member of society. She either cannot get knowledge by believing on the basis of this testimony or cannot get moral credit for acting on the basis of such testimony.[11] In short, this strong moral epistemic component neither seems to be the special province of philosophers nor seems to be conducive to showing what philosophers can do for social moral belief formation.

One way of weakening Strong Moral Expertise, however, is precisely to relax the normative epistemic component in the tripartite structure as follows for *Weak Moral Experts*:

Definition 2 *(Weak Moral Experts). Those who know some set of normative positions which are taken to be potentially correct. Furthermore, they are able to (at least reliably) validly infer morally relevant and important conclusions from these normative position and other relevant non-normative facts.*

By weakening the normative epistemic component from knowledge that the correct normative position is true to knowledge *of* some set of normative positions which are taken to be potentially correct, this definition of Weak Moral Experts starts to look more like the domain of philosophers. The normative positions taken to be

[10] The issue of determining expertise arises also in multiple forms throughout Ericsson et al. (2006). Goldman (2001) calls this the expert/expert problem. However, note that, since on my definitions, Weak Moral Expertise is always a matter of degree and comparison classes, there is no hard and fast distinction between novice/expert problems and expert/expert problems as Goldman distinguishes them.

[11] Others have made distinct objections that apply to Strong Moral Expertise. While I am agnostic about these, one sample objection comes from Archard (2011), who suggests that moral philosophy is built on—or heavily draws on—common-sense morality. If so, it may be questionable for philosophers to view themselves as capable of moral judgment inaccessible to members of the public.

potentially correct could be positions that are currently considered live possibilities amongst the philosophical community.[12] Minimal requirements for being potentially correct include consistency, but I wish to leave open the question of which additional theoretical virtues to adopt.[13]

On this definition, Weak Moral Expertise is a weaker notion than Strong Moral Expertise in the sense that all Strong Moral Experts are also Weak Moral Experts, but not vice versa. This is because a Weak Moral Expert should be aware of *potentially* correct normative positions as well as the relevant non-normative facts, but does not know *which* normative position is correct.[14] There is another important difference between Strong and Weak Moral Expertise. While Strong Moral Experts know that their inferred conclusions have the virtue of truth, Weak Moral Experts are not able to determine whether their valid inferences yield truth. However, they are more Weakly Moral Expert when they are able to judge or evaluate the importance and novelty of these inferences.[15] Novelty has to do with the straightforward question of whether the inferences have been made before; arguments are more important or interesting when they interact with our (either philosophical or folk) intuitions in surprising ways.[16]

[12] Alternatively, we could take the potentially correct normative positions to be those that are "reasonable comprehensive doctrines" in society, using the Rawlsian terminology (Hegstad, 2023). However, this seems less satisfactory to me, since the philosopher is better placed to understand which moral theories are philosophically respectable than which are socially respectable.

[13] I follow the judgments of O'Hear (2009) both that in "the choice of premises the will is free: logic cannot dictate the ground from which its conclusions proceed" (p. 353) and that (at least) with respect to the practical "there can be many starting points and many goals which are not simultaneously reconcilable" (p. 357).

[14] Although I say it is weaker, that's a practical claim. It is logically possible that one could be a Strong Moral Expert in the sense of knowing or reliably adopting (e.g. by perfect intuition assuming some realist ontology) the correct normative position, but while having no conception of competing potentially correct positions. Such an individual would be a Strong Moral Expert but not a Weak Moral Expert in my sense. However, this seems to be a case which can be set aside as unlikely to arise; it is difficult to see how an individual could come to adopt the correct normative position without considering alternative normative positions and comparing them. To be provocative, a moral prodigy is harder to imagine than, say, a chess or violin prodigy.

[15] As a nice corollary, adopting Weak Moral Expertise helps retain a notion of moral philosophical progress, which Strong Moral Expertise does not. As we better understand what follows from various normative positions, we have the opportunity—which some may take and others reject—of revising them to see how to avoid counterexamples and rebuttals but more structurally to understand what they imply, both in theory and in practice. Importantly, this understanding neither depends on nor predicts the claim that the substantive (non-conditional) commitments will converge. However, greater understanding is an important form of progress. In contrast, there is no guarantee that more people will be Strong Moral Experts or that Strong Moral Expertise grows over time. (Presumably, just as Strong Moral Experts argue (soundly) for the moral truth, other convincing rationales can be given for moral falsehoods—it is not clear that one must win, especially given the similar levels of support for incompatible philosophical positions (Chalmers, 2015).)

[16] In private correspondence, Michael Rabenberg points out an interesting (conservative) implication of this account. One can be a Weak Moral Expert purely by being fully up-to-date with moral philosophy. This yields a vast amount of inferences at your fingertips, the capacity to judge whether either (a) those inferences are novel (by stipulation, none are) or (b) those inferences are important or interesting (some will be and some will not). Of course, this implication reinforces the point that moral philosophers will tend to be well-placed to play the role of Weak Moral Experts.

These definitions also draw attention to the essentially *comparative* nature of expertise.[17] Knowledge of more normative positions, knowledge of more (relevant) non-normative facts, and greater ability to deduce sound conclusions all serve to increase one's Strong Moral Expertise. Familiarity with more potentially correct normative positions, knowledge of more (relevant) non-moral facts, and greater facility with drawing valid inferences all serve to increase one's Weak Moral Expertise. Similarly, greater evaluative ability to recognize novelty and importance increase one's Weak Moral Expertise. One can adapt either definition to require some threshold for any of accounts' components, but I can be agnostic about whether to adopt such thresholds here.

Drawing on the comparative nature of expertise, one question that arises is whether Weak Moral Experts should privilege *deepening* their understanding of specific (perhaps endorsed) normative positions (acting in a "partisan role" (Mintz-Woo, 2025)) or should privilege *widening* their understanding to a range of normative positions (acting as a "conduit for the discipline" by channeling a variety of normative positions from the philosophical literature). As noted above, both strategies can increase Weak Moral Expertise. I think the answer is that the answer should be sensitive both to one's expected role and, crucially, to one's expectations about *others'* behavior. For instance, if one expects to be helping specific particular people reach justified moral judgments (e.g., as a clinical bioethical specialist), then it is valuable to have familiarity with a wide variety of normative positions (i.e., widening or being a conduit for the discipline is an appropriate strategy). This is so one can assist more patients who may endorse a wide variety of values. In contrast, if one is translating philosophy to the public in newspaper opinions, generally one is explaining some argument or inference, so it is less important that one is familiar with a variety of normative positions (i.e., deepening or acting as a partisan seems more appropriate).

However, deepening in order to interact at the philosophy-public interface is justified when one also expects that others will be able to provide contrasting arguments from distinct normative positions. In case one expects few others to contribute such arguments, widening makes sense so that the public can get a sense of the normative debate. For instance, it is permissible to explain the utilitarian argument for being vegetarian if one expects that there are others who will take distinct normative positions and explain their implications—whether similar or different. If one expects that there are no others, it is more important that one also tries to indicate how other normative positions bear on the moral question. As above, a clinical bioethicist is expected to (alone) help patients come to moral judgments; it is important that they are able to work with a variety of normative positions and principles and explain how they are relevant, since, within some reasonable range, they should be able to work with the values of the patient herself.

This parallels a debate in the ethics of policy commissions. In this debate, Hegstad (2023) argues that, insofar as one expects that moral experts representing a variety of moral views will be on a commission, it is appropriate for any given expert to defend her own chosen normative position (also cf. Hegstad, 2024). In contrast, if one

[17] Note that the comparative nature of moral expertise adduced here helps address the concerns of Priaulx et al. (2016) that ascribing moral expertise need imply *absolute* authority.

expects that few others will contribute philosophy-public arguments and inferences, it might be appropriate to widen and indicate how some variety of moral positions are relevant. This parallels an argument I make elsewhere (2025) when discussing the case of policy commissions with a lone (or almost lone) philosopher. In such a case, one might want to engage in a "convergent evaluation" and take the role of a "conduit" for philosophy, whereby the philosopher lays out a variety of moral positions and evaluates where they converge or diverge, instead of advocating her own ("partisan") position, which is like narrowing.

Having explained these different responses to moral expertise, I now turn to why moral philosophers are most appropriately thought of as experts of the Weak Moral Expert sort. I do think that philosophers are trained to focus on the skilled component (drawing valid inferences) more explicitly and more carefully than in most other disciplines, but this could be controversial.[18] What should not be controversial is that the *normative epistemic component* of Weak Moral Expertise, i.e., knowledge of potentially correct normative positions and their workings, *is* to a significant extent the special province of philosophers. Familiarity with these positions makes it easier to apply them and to recognize when proffered moral inferences are novel. In other words, philosophical expertise (at least in the context of moral domains) is generally held by those with comparatively high levels of Weak Moral Expertise, that is, moral philosophers.[19] For those who would claim that it is helpful to improve the reasoning and belief-formation *processes* of the public, e.g., by drawing attention to heuristics and biases, I agree—but I disagree that this is the role of the philosopher. Generally speaking, improving public reasoning is better suited to those familiar with relevant empirical literature, such as psychologists. Obviously, some philosophers also have this kind of expertise, but I would suggest that it is further from our core competences.

Finally, it is worth pointing out the importance of the non-normative facts. These are sometimes neglected in discussions of moral expertise but, without them, normative positions do not ground practical judgments.[20] One thing to note is that knowledge of relevant non-normative facts is often highly domain-dependent, which

[18] Confer Prinzing and Vazquez (2025), who find that studying philosophy predicts stronger reasoning skills and development of some intellectual virtues. This builds on previous work by Prinzing and Vazquez (2024), who canvassed the empirical literature and found mixed and unclear evidence (in particular, that the empirical literature could not distinguish treatment from selection effects with respect to studying philosophy). But note that *all* the empirical evidence canvassed by Prinzing and Vazquez (2024) is grist for my mill; philosophers who selected into the discipline and are strong reasoners are just as valuable for my claims as philosophers who developed reasoning skills *as a result of* their philosophical training.

[19] It is worth noting that this is not to deny that philosophers, collectively and individually, are subject to heuristics and biases. Even worse, these biases may apply even to topics that are supposed to be in areas of philosophical expertise (Schwitzgebel & Cushman, 2012, 2015). However, I believe that these biases relate more to substantive judgments, if indeed they apply at all, as opposed to reasoning skills and capacity to develop—or check the validity of—inferences and arguments. As noted above, I believe that Prinzing and Vazquez (2024, 2025) have made a solid case that those who have studied philosophy tend to have stronger skills related to inferences and argumentation.

[20] There are exceptions. McGrath (2009), for instance, notes the importance of factual information in the context of moral expertise by distinguishing between "pure" and "impure" forms of moral deference. However, even this terminology, I would submit, reflects a privileging of the moral and a downgrading of the non-moral.

correctly implies that, since few of us have time to master many factual domains, most of us can only aspire to become moral experts in, at best, a limited number of them.

What is the purpose of this section? If philosophers have the kind of expertise required to be Weak Moral Experts, and Weak Moral Expertise does not face the same kind of objections as Strong Moral Expertise, then there is a less- or un-controversial claim that philosophers have this kind of expertise (Chick & LaVine, 2014). The importance of this is that it makes clear exactly what it is that philosophers can *generally* do better than non-philosophers, contra the skeptics. Adopting Weak Moral Expertise also lays out the groundwork for what moral philosophers can contribute to society.

3 What can moral philosophers contribute to society?

De Shalit (2020) expands the familiar Rawlsian notion of reflective equilibrium to incorporate the public at large. "Reflective equilibrium," initially outlined by Rawls (1951) (albeit without that moniker), is the term for a method of moral adjudication which weighs considered practical moral judgments with mid- and high-level principles, adjusting as necessary until the components cohere and explain each other.

De Shalit suggests that the process should not be performed by the lone philosopher, but should also involve the public. He has a couple reasons for this: firstly, it is more consistent with political philosophers' professed focus on democratic institutions and theories; secondly, it could be more fruitful and creative. I endorse this appeal to heterogeneous and diverse public values, but will argue that his formulation is not the best way of respecting them. His proposed expansion we can call "public initial reflective equilibrium", by appealing to the distinction[21] between initial (hypothesis generation) and final (evidence evaluation and conclusion drawing) stages:[22]

Definition 3 *(Public Initial Reflective Equilibrium). The version of reflective equilibrium where public values or intuitions are introduced as inputs into a philosopher's reflective equilibrium process.*

[21] This distinction is meant to evoke a classic one from philosophy of science: the "context of discovery" and the "context of justification" (Mintz-Woo, 2020). This "context distinction" was introduced by Hans Reichenbach of the Berlin Circle to indicate parts of the (idealized) scientific process. Roughly speaking, the context of discovery involves investigating phenomena, generating hypotheses and gathering evidence while the context of justification involves the logical processes of linking the evidence to conclusions or the construction of arguments for substantive conclusions. While there are both objections to the distinction (Nickles, 1985) and controversies about how to understand it (Schickore & Steinle, 2006), it has provided a basis for discussion in philosophy of science throughout much of the twentieth century (Hoyningen-Huene, 1987).

[22] He calls his view "public reflective equilibrium" which contrasts with my "public final reflective equilibrium". To distinguish his method from mine, I call his "public initial reflective equilibrium." His initial use of "public reflective equilibrium" was in order to contrast it with what he calls Rawls' theory, "private reflective equilibrium". Sadly, in nomenclature, ungainliness begets ungainliness. I regret playing my part and welcome suggestions on improved terminology.

The method is to bring the values of the public in at the ground level, for instance, by having philosophers randomly discuss (political) issues with people in society and incorporate their suggested moral and political judgments and principles into the philosophers' moral theorizing (de Shalit, 2020). To illustrate, the philosopher could engage random citizens on the bus and ask them about their views on immigration. After hearing some range of judgments, proto-principles and arguments, the philosopher takes some of this as inputs into her reflective equilibrium process (this may well involve pruning or systematizing the public input). This is public initial reflective equilibrium because the public values or intuitions come in initially as input *into* the philosopher's reflective equilibrium process. Weighing and considering her own judgments and theories, and including the public values, the philosopher herself engages in a reflective equilibrium process.

But I believe that gathering the views of the public does not play to the strengths of the philosopher (or not, at least, to many philosophers)—plausibly, gathering the views of the public is the expertise of social scientists (e.g., demographers, sociologists, or social psychologists) (Hegstad, 2023; Mintz-Woo, 2025). The strength of the philosopher lies elsewhere, as argued above, in validly inferring novel conclusions from normative positions.

It is also an issue since, while some people are very engaged in policy and politics, it is entirely reasonable for people to profess few or no (systematic or articulable) moral commitments. So we might have a lot of trouble eliciting such initial public values. While we might wish it were the case that people reflected more on their values, holding or discussing substantive normative positions is not required of any citizen.

For what I call "public final reflective equilibrium", in contrast to de Shalit's model, the philosophers work out potential theories as best they can in light of their Weak Moral Expertise, explain the links between those theories and practice, and then the resulting normative positions face the tribunal of social values not individually but only as a corporate body:[23]

Definition 4 *(Public Final Reflective Equilibrium). The version of reflective equilibrium where philosophers' arguments or inferences are used as inputs into the reflective equilibrium process(es) of (members of) the public, who evaluate normative positions based on their values, considered judgments, and the philosophers' arguments or inferences.*

The idea here is that the worked-out implications of normative positions, in the form of valid inferences and arguments, are communicated or made available to the public (e.g., in newspapers, magazines or online repositories), who can, at that point, compare and contrast different normative positions (Pigliucci & Finkelman, 2014, have a good list of examples). In a word, Weak Moral Experts contribute to greater public *understanding* of normative positions, i.e., comprehension of the links and connections between different claims and positions (Hannon & Nguyen, 2022). The goal of this method is for members of society to adopt normative positions which

[23] With apologies to Quine (1951).

reflect their (reasonably) individually inflected values and commitments (similar ideas are expressed by O'Hear, 2009). Again, the main point of contrast is as follows: instead of eliciting public values and then having the philosopher engage in reflective equilibrium with those public values incorporated (initial public reflective equilibrium), the philosopher provides the arguments and (members of) the public engage(s) in processes of reflective equilibrium (final public reflective equilibrium). These processes might generate different outputs. Even if we have established which implications follow from a set of theories, that need not determine how to select amongst them (here, I concur deeply with Beebee, 2018). Different individuals might well weigh what I would call these "baskets of theories and implications" differently.

To illustrate, the picture is as follows: suppose some person is trying to come to a practical conclusion in a morally-laden context. She has some values and some moral intuitions, but isn't clear about how these bear on the context in question. Public final reflective equilibrium is the view that she would benefit from exposure to philosophical arguments, arguments which link values and moral premises to practical or applied conclusions. So, for instance, public philosophy (e.g., via radio programs or television articles) explaining how virtue-formation is relevant to anti-racism, or how consideration of pleasure has implications for (vegetarian) diets, or how rights support climate action could help her come to connect her values with various practical concerns (cf. Chick & LaVine, 2014). After this philosophical input, she engages in reflective equilibrium, considering how various baskets of theories plus consequences compare to other baskets, leading her to find some moral theories more acceptable than others. This is final reflective equilibrium because the philosophical work is done and communicated to the member of the public and that person then evaluates the arguments and implications, weighing them in a personal reflective equilibrium. This contrasts with public initial reflective equilibrium, where the public enters into the initial phase, e.g., by providing their normative intuitions or judgments as inputs into the philosopher's reflective equilibrium.

It is worth clarifying that this is meant to answer the question of how moral philosophers can help members of the public come to justified or informed moral beliefs. We can say that this question occurs at the "philosophy-public" interface. It is not—at least not directly—about how to aggregate social views or generate public deliberative bodies or other such social or collective belief-forming structures aimed at policy formation (for such processes, cf. Christiano, 2012). We could say these other procedures are at the "philosophy-policy" interfaces. Final public reflective equilibrium is "public" in the sense of relating to the public, not in terms of being some type of public forum or deliberation procedure—and certainly not in terms of being intended (only) for public policy or governance.

Furthermore, the role of philosophers I am suggesting is to provide relevant input: how normative positions have implications for practical judgments. The role is not to prescribe particular practical judgments; there might be a range of reasonable disagreement with respect to the weight put on given values and principles, as well as with given practical judgments. Given this range, one might hold that adopting specific practical ("expert") advice is incompatible with important forms of autonomy (D'Agostino, 1998). Since philosophers are not providing the moral values, but just linking particular normative positions with particular practical judgments via valid

inference, the approach defended here does not undermine autonomy for individual moral belief-formation.

There are several preconditions for the success of public final reflective equilibrium. First, it clearly requires that some philosophers make their work comprehensible to society.

A closely related precondition is that the arguments and texts are themselves accessible to the public. Since many philosophical works are behind paywalls or expensive to access, this is in no way trivial. However, many philosophers do engage with the public, either by writing popular discussions of philosophical topics or by applying philosophical theories to contemporary issues (a handful of recent exemplars include Yancy, 2015; Manne, 2018; Nguyen, 2018; Rini, 2020; Táíwò & Cibralic, 2020). This requires extra effort—and public philosophy is not part of the standard philosophical skill set—but it could open up new scholarship ideas and opportunities (Chick & LaVine, 2014). At this point, this is done in a piecemeal fashion, so there are fewer opportunities to holistically compare the arguments and claims presented publicly to their alternatives. That, of course, is a theoretical justification for doing *more* public philosophy, not less. This contention is also supported by the fact that most philosophers' salaries are paid for by the public, meaning that delivering our inferential results in ways that are accessible to the public might be thought of as a minimal obligation in return.

Another precondition is that members of society are interested in engaging with moral theory and argumentation. We are considering a particular member of society who wants to determine how she should form or choose moral beliefs, so she is likely to be motivated in this way, but she may or may not be representative. Here, we can expect that some politically salient moral topics, such as those involved with food ethics, climate ethics, social philosophy and ways of responding to various kinds of structural injustices, might find receptive audiences. However, many others (especially recondite discussions like those in metaethics) might not. I do not think that there is any epistemic requirement that members of society participate in or reflect on philosophical discussion, but if they are presented with the practical import of well-understood normative positions, some will engage. I also do not think there is any professional duty for any given philosopher to engage in this manner either. But it would be good for some philosophers to engage in this way, and those who do should be commended for doing so. Of course, there is also the non-trivial co-benefit of making philosophy less mysterious and more recognizable (Pigliucci & Finkelman, 2014).

However, if these preconditions can be met, what is the theoretical case for adopting public final reflective equilibrium over public initial reflective equilibrium? The most important objection to public initial reflective equilibrium is that the public values end up having less weight, both because they are more easily rejected by the philosopher and because they are uninformed by the philosophical inferences that might help strengthen them. In other words, bringing them in as inputs artificially limits their role in the reflective equilibrium process in two ways: first, via the potential rejection of public values by the philosopher in the reflective equilibrium stage; second, via the inability of the public to give comparative and holistic evaluations of the normative positions. I will expand on these in turn.

With public initial reflective equilibrium, inputs can be lost or jettisoned in the reflective stage performed by the philosopher, whereas public final reflective equilibrium involves the public in the moral context of evaluation. Why is this? If the interlocutors' ideas are too inchoate or implausible *as judged by the philosopher herself*, then they do not make it into the reflective equilibrium at all. One might not think this is a problem; the philosopher is after all the relevant professional in this exchange. However, if we have rejected the Strong Moral Expert role for the philosopher, then we do not think that being the relevant professional is enough for her substantive value judgments to be reliably more accurate. A closely related worry is that the interlocutors might be unable to express their view successfully, or in terms that the philosopher can follow, perhaps leading to these views being dismissed due to phrasing instead of due to plausibility or truth-value.

The second reason that it is less theoretically valuable for public values to be brought in at the first stage is that the public is not able to holistically compare and contrast competing normative theories *until* their implications have been well-understood—the comparative advantage, I have argued, of the philosopher. This is because the philosopher can fill in those baskets of theories and implications so the costs and benefits of accepting a basket are clearer. If the public is unaware of the range of potentially correct normative positions with respect to some given subject matter, they will not properly evaluate the position in comparison to others. This could, inter alia, lead to costs being overly emphasized. For instance, suppose a normative position is committed to some unintuitive judgment but all or most competing positions turn out to be committed to the *same* judgment. If considering that position in isolation, one might not recognize this convergence while weighing whether or not to adopt the theory (situations like this are sometimes referred to as "companions in guilt" since all the normative positions end up being "guilty" together).

Finally, before philosophers have worked out greater understanding of particular normative positions, it may be easy to focus on some aspects without recognizing their broader implications (which of course may be even more damaging or exonerating). If not presented with the strongest, most worked-out version of a normative position, one could dismiss it prematurely. Since it is, qua Weak Moral Expert, a role that philosophers can play: to work out all the valid (novel, important) implications of their positions, this is a way that the philosophical (or, more precisely, the Weak Moral Expert) community could help the public in moral deliberations.

I take these to be decisive reasons to favor public final reflective equilibrium over public initial reflective equilibrium. While public final reflective equilibrium may require more philosophical outreach, I believe that the resulting picture not only provides a specifically appropriate role for the moral philosopher, it also provides the opportunity for society to understand and engage with philosophers, something that is also potentially valuable for the sustainability of the philosophical discipline as a whole.

4 Objections

In this section, I consider two objections. The first has to do with potential divergences between philosophers and the public. The second is rooted in the concern that Strong Moral Expertise is phenomenologically privileged.

First, consider an objection about the role of philosophers: what if the views shared by philosophers differ from the views shared by the public (or by some salient subpopulations)? Given the professed importance of public values and the arguments for their incorporation into public final reflective equilibrium, we might be concerned if there is a radical mismatch.[24] This might seem to support public initial reflective equilibrium—as philosophers are confronted with different types of people with different types of values, wouldn't this act as a corrective on the types of values held by philosophers? After all, we know that professional philosophers are not representative of the population as a whole (Botts et al., 2014). Why should we expect that their values will be representative?

I want to grant my interlocutors the possibility—indeed, the likelihood—that moral philosophers may have values that do not represent the public.[25] Nevertheless, I believe public final reflective equilibrium is preferable to public initial reflective equilibrium, since working out the consequences of theories that others disagree with in greater detail does not ipso facto make those theories more plausible or more acceptable once the theories and their implications are made available. Indeed, I would hazard the guess that working these consequences out in greater detail makes them more risible to those who disagree. (I grant that this is ultimately an empirical supposition, and one I would be pleased to see confirmed—or disconfirmed if incorrect.)

In other words, it is not necessary for this picture that philosophers have the same values as the public: they just work out the implications of various values that they are interested in (whether because they endorse or deny those values). Members of the public can consider these arguments and take them to be reductios or constructive results depending on their own personal positions. To the extent that a given member of the public has values that are not considered by philosophers at all, philosophy will have less to offer that individual (not a surprising conclusion). At the very least, that person can consider extant arguments as reductios or consider them as ways to reject other values (and support her own). It is also possible that she could try to convince philosophers to consider her values or, if she were so inclined, develop her own moral theories. After all, finding some moral positions that are undefended or

[24] Thanks to Liam Kofi Bright for bringing this objection up in personal discussion.

[25] Interestingly, this is a place where the limited scope of my claims to moral or normative claims differs from more broad philosophical progress. While the views of other philosophers—metaphysicians, for instance—likely differ from those of the public on metaphysical theories, we can expect that (a) members of the public have unconsidered views on metaphysical issues and (b) insofar as they do, those views would be more lightly held. Moral and political claims look different; the public is more likely both to have considered these claims and to have their moral views reflected in their identities or their everyday lives. In other words, we can expect that this objection would be *less* of a problem with *non-moral* claims of philosophical expertise.

underconsidered is quite possibly the route to a philosophical career—and quite challenging to do.

The second objection is about how the philosophers experience their roles. Some philosophers could object that their role is not merely to generate new conditional claims, or that their work would be unsatisfying or insufficiently important if they were not seeking absolute truth as opposed to merely conditional truth. They might say that what they take themselves to be doing is more in line with Strong Moral Expertise than Weak Moral Expertise, and we should take this phenomenology seriously.[26]

This can be addressed more straightforwardly. There is a debunking explanation with respect to the philosopher's usual role, drawing valid implications. From the point of view of a Weak Moral Expert who is working with the normative positions she *favors or believes to be correct*, the advancing of the favored position is phenomenologically identical to the Strong Moral Expert identifying valid arguments connecting the *true* moral theory to new and interesting conclusions. So a philosopher may take herself to be a Strong Moral Expert, but if she was mistaken about this (and merely though herself to be a Strong Moral Expert), it would not be phenomenologically distinctive enough for her to recognize this.

5 Conclusion

It may seem unsurprising that a moral philosopher would counsel a member of the public trying to come to true moral beliefs to listen to moral philosophers. However, this paper has argued for this claim in a way that does not depend upon or attribute to philosophers reliable moral truth-tracking. In other words, the goal was to show that moral philosophers are useful given fairly minimal assumptions about the skills and capacities of moral philosophers. They can help society by facilitating public *understanding* of normative positions.

In so doing, I hope to have made the case the moral philosophers do have a form of moral and philosophical expertise, but one which does not require many controversial commitments. Moral philosophers can take the role of Weak Moral Experts not only because this gives some guidance about what they can produce, but also because may be better placed than other members of society to play this role.

In turn, Weak Moral Experts can advance individual moral reasoning by discovering and considering new implications of and valid arguments for various moral theories. Communicating these arguments through public media, these Weak Moral Experts can help remove the mystery of what philosophers do as well as promote understanding of moral theory. These arguments would help individuals both decide on practical judgments as well as evaluate different moral theories by examining the implications of these theories, aiding moral belief formation. They can do this whether or not they substantively endorse a given theory.

[26] Hursthouse (1991) for instance, makes a related objection from the point of view of someone *seeking* moral expertise. She asks whether it would be sensible to ask for advice from one who adverted to uncertainty about the true moral theory; she thinks this evidently absurd (232).

Finally, this picture can be expanded to incorporate society as a complement to the Weak Moral Experts since, given the rejection of Strong Moral Expertise, members of society may be as well placed as anyone else to develop moral beliefs. As the baskets of theories plus implications are more comprehensively explored, the strengths and weaknesses of the various theories can be better evaluated accordingly (with different weightings of those strengths and weaknesses explaining why this may not lead to descriptive convergence—or even normative convergence (Mintz-Woo, 2025)).

With all of these components together, even given concerns about moral methodology and moral testimony—which may well be pessimistic—we can still allow for philosophers to contribute to society's belief formation as we expand our understanding of the implications of our moral beliefs.

Acknowledgements The ideas in this paper have been in development for a very long time (the first formal version of it was presented at the Society for Applied Philosophy's Annual Conference 2017 and the most recent version at CEPPA at the University of St Andrews in 2025), and there are many people who have contributed in ways small and large, but I especially want to thank Liam Kofi Bright, Michael Hannon, Philip Kitcher, Elizabeth O'Neill, Tom Parr, Avia Pasternak, Michael Rabenberg, Don Ross, and Peter Singer. Many thanks to the reviewers and to the editors, and in particular to the guest editors, especially Cuizhu (Dawn) Wang. Some of the ideas here were informed by a volunteer expert position in the Irish Carbon Budgets Working Group (2024–2025).

Funding Open access funding provided by International Institute for Applied Systems Analysis (IIASA). Open access funding for this article comes from KËMO, the Austrian consortium for research libraries.

Declarations

Conflicts of interest The author has volunteered in assisting national public policymaking and is interested in philosophers making productive contributions to public processes. The author declares no other conflicts of interest.

References

Archard, D. (2011). Why moral philosophers are not and should not be moral experts. *Bioethics*, *25*(3), 119–127. Retrieved from https://doi.org/10.1111/j.1467-8519.2009.01748.x

Baz, A. (2016). Ordinary language philosophy. In H. Cappelen, T. S. Gendler, & J. Hawthorne (Eds.), *The Oxford handbook of philosophical methodology* (pp. p. 112—-129). Oxford University Press. Retrieved from https://doi.org/10.1093/oxfordhb/9780199668779.013.28

Beebee, H. (2018). Philosophical scepticism and the aims of philosophy. *Proceedings of the Aristotelian Society*, *118*(1), 1–24, Retrieved from https://academic.oup.com/aristotelian/article/118/1/1/4924516

Birnbacher, D. (2012). Can there be such a thing as ethical expertise? *Analyse & Kritik*, *34*(2), 237–249, Retrieved from http://analyse-und-kritik.net/2012-2/AKBirnbacher2012.pdf

Botts, T. F., Bright, L. K., Cherry, M., Mallarangeng, G., & Spencer, Q. (2014). What is the state of Blacks in philosophy? *Critical Philosophy of Race*, *2*(2), 224–242. Retrieved from https://doi.org/10.5325/critphilrace.2.2.0224

Bourget, D., & Chalmers, D. J. (2014). What do philosophers believe? *Philosophical Studies*, *170*(3), 465–500. Retrieved from https://doi.org/10.1007/s11098-013-0259-7

Bourget, D., & Chalmers, D. J. (2023). Philosophers on philosophy: The 2020 PhilPapers survey. *Philosophers' Imprint*, *23*, 11. Retrieved from https://doi.org/10.3998/phimp.2109

Chalmers, D. J. (2015). Why isn't there more progress in philosophy? *Philosophy*, *90*(1), 3–31. Retrieved from https://doi.org/10.1017/S0031819114000436

Chick, M., & LaVine, M. (2014). The relevance of analytic philosophy to personal, public, and democratic life. *Essays in Philosophy*, *15*(1), 138–155. Retrieved from https://doi.org/10.7710/1526-0569.1494

Cholbi, M. (2007). Moral expertise and the credentials problem. *Ethical Theory and Moral Practice*, *10*(4), 323–334. Retrieved from https://doi.org/10.1007/s10677-007-9071-9

Christiano, T. (2012). Rational deliberation among experts and citizens. In J. Parkinson & J. Mansbridge (Eds.), *Deliberative systems* (pp. 27–51). Cambridge University Press. Retrieved from https://doi.org/10.1017/CBO9781139178914.003

Cosker-Rowland, R. (2017). The epistemology of moral disagreement. *Philosophy Compass*, *12*(2), e12398. Retrieved from https://doi.org/10.1111/phc3.12398

Crosthwaite, J. (1995). Moral expertise: A problem in the professional ethics of professional ethicists. *Bioethics*, *9*(5), 361–379. Retrieved from https://doi.org/10.1111/j.1467-8519.1995.tb00312.x

D'Agostino, F. (1998). Expertise, democracy, and applied ethics. *Journal of Applied Philosophy*, *15*(1), 49–55. Retrieved from https://doi.org/10.1111/1468-5930.00072

de Shalit, A. (2020). Political philosophy and what people think. *Australasian Philosophical Review*, *4*(1), 4–22. Retrieved from https://doi.org/10.1080/24740500.2021.1876411

Doris, J. M. (2021). *Charater trouble*. Oxford University Press. Retrieved from https://doi.org/10.1093/oso/9780198719601.001.0001

Ericsson, K. A., Charness, N., Feltovich, P. J., & Hoffman, R. R. (2006, June). *The Cambridge handbook of expertise and expert performance*. Cambridge University Press. Retrieved from https://doi.org/10.1017/CBO9780511816796

Gesang, B. (2010, May). Are moral philosophers moral experts? *Bioethics*, *24*(4), 153–159. Retrieved from https://doi.org/10.1111/j.1467-8519.2008.00691.x

Goldman, A. I. (2001). Experts: Which ones should you trust? *Philosophy and Phenomenological Research*, *63*(1), 85–110. Retrieved from https://doi.org/10.1111/j.1933-1592.2001.tb00093.x

Hannon, M., & Nguyen, J. (2022). Understanding philosophy. *Inquiry, Inquiry*. Retrieved from https://doi.org/10.1080/0020174X.2022.2146186

Hegstad, E. (2023). Moral experts as members of ethics commissions as seen through the prism of comprehensive doctrines. *Bioethics*, *37*(6), 543–550. Retrieved from https://doi.org/10.1111/bioe.13086

Hegstad, E. (2024). Why ethics commissions? Four normative models. *Res Publica*, *30*(1), 67–85. Retrieved from https://doi.org/10.1007/s11158-023-09622-2

Hills, A. (2009). Moral testimony and moral epistemology. *Ethics*, *120*(1), 94–127. Retrieved from https://doi.org/10.1086/648610

Hills, A. (2013). Moral testimony. *Philosophy Compass*, *8*(6), 552–559. Retrieved from https://doi.org/10.1111/phc3.12040

Horvath, J., & Koch, S. (2020). Experimental philosophy and the method of cases. *Philosophy Compass*, *16*(1), e12716. Retrieved from https://doi.org/10.1111/phc3.12716

Hoyningen-Huene, P. (1987). Context of discovery and context of justification. *Studies in History and Philosophy of Science*, *18*(4), 501–515, Retrieved from https://doi.org/10.1016/0039-3681(87)90005-7

Hursthouse, R. (1991). Virtue theory and abortion. *Philosophy and Public Affairs*, *20*(3), 223–246, Retrieved from http://www.jstor.org/stable/2265432

Knobe, J. (2003). Intentional action and side effects in ordinary language. *Analysis*, *63*(3), 190–194. Retrieved from https://doi.org/10.1111/1467-8284.00419

Knobe, J. (2007). Experimental philosophy and philosophical significance. *Philosophical Explorations*, *10*(2), 119–121. Retrieved from https://doi.org/10.1080/13869790701305905

Manne, K. (2018). *Down girl: The logic of misogyny*. Oxford University Press. Retrieved from https://doi.org/10.1093/oso/9780190604981.001.0001

McGrath, S. (2009). The puzzle of pure moral deference. *Philosophical Perspectives*, *23*(1), 321–344. Retrieved from https://doi.org/10.1111/j.1520-8583.2009.00174.x
McIntyre, L. (2022). What is public philosophy? In L. McIntyre, N. McHugh, & I. Olasov (Eds.), *A companion to public philosophy* (pp. 3–8). Wiley. Retrieved from https://doi.org/10.1002/9781119635253
Mintz-Woo, K. (2020). Public values in the right context. *Australasian Philosophical Review*, *4*(1), 57–62. Retrieved from https://doi.org/10.1080/24740500.2020.1926891
Mintz-Woo, K. (2025). Explicit methodologies for normative evaluation in public policy, as applied to carbon budgets. *Journal of Applied Philosophy*. Retrieved from https://doi.org/10.1111/japp.70047
Nguyen, C. T. (2018). Escape the echo chamber. *Aeon*, Retrieved from https://aeon.co/essays/why-its-as-hard-to-escape-an-echo-chamber-as-itis-to-flee-a-cult
Nickel, J. W. (1988). Philosophy and policy. In D. M. Rosenthal & F. Shehadi (Eds.), *Applied ethics and ethical Theory, chapter 7* (pp. 139–148). University of Utah Press.
Nickles, T. (1985). Beyond divorce—current status of the discovery debate. *Philosophy of Science*, *52*(2), 177–206. Retrieved from https://doi.org/10.1086/289239
Nielsen, K. (1987). Can there Be progress in philosophy? *Metaphilosophy*, *18*(1), 1–30. Retrieved from https://doi.org/10.1111/j.1467-9973.1987.tb00184.x
O'Hear, A. (2009). Philosophy – wisdom or technique? *Royal Institute of Philosophy Supplements*. *65*, 351–361. Retrieved from https://doi.org/10.1017/S1358246109990166
Peirce, C. S. (1931–1958). *Collected papers of Charles Sanders Peirce*. Harvard University Press.
Pigliucci, M., & Finkelman, L. (2014). The value of public philosophy to philosophers. *Essays in Philosophy*, *15*(1), 86–102. Retrieved from https://doi.org/10.7710/1526-0569.1491
Priaulx, N., Weinel, M., & Wrigley, A. (2016). Rethinking moral expertise. *Health Care Analysis*, *24*(4), 393–406. Retrieved from https://doi.org/10.1007/s10728-014-0282-7
Prinzing, M., & Vazquez, M. (2024). Does studying philosophy make people better thinkers? *Journal of the American Philosophical Association*, *10*(4), 855–876. Retrieved from https://doi.org/10.1017/apa.2023.30
Prinzing, M., & Vazquez, M. (2025). Studying philosophy does make people better thinkers. *Journal of the American Philosophical Association*, *24*(4), 393–406. Retrieved from https://doi.org/10.1017/apa.2025.10007
Quine, W. V. O. (1951). Main trends in recent philosophy: Two dogmas of empiricism. *The Philosophical Review*, *60*(1), 20–43. Retrieved from https://doi.org/10.2307/2181906
Rasmussen, L. M. (2005). Introduction: In search of ethics expertise. In L. M. Rasmussen (Ed.), *Ethics expertise: History, contemporary perspectives, and Applications, chapter 1* (pp. 1–14). Springer. Retrieved from https://link.springer.com/chapter/10.1007/1-4020-3820-8_1
Rawls, J. (1951, April). Outline of a decision procedure for ethics. *The Philosophical Review*, *60*(2), 177–197, Retrieved from http://www.jstor.org/stable/2181696
Rini, R. A. (2020). The internet is an angry and capricious god. *Times Literary Supplement*, , Retrieved from https://www.the-tls.co.uk/articles/the-internetis-an-angry-and-capricious-god/
Schickore, J., & Steinle, F. (Eds.). (2006, June). *Revisiting discovery and justification: Historical and philosophical perspectives on the context distinction*. Dordrecht: Springer. Retrieved from https://doi.org/10.1007/1-4020-4251-5
Schwitzgebel, E., & Cushman, F. (2012). Expertise in moral reasoning? Order effects on moral judgment in professional philosophers and non-philosophers. *Mind & Language*, *27*(2), 135–153,Retrieved from https://doi.org/10.1111/j.1468-0017.2012.01438.x
Schwitzgebel, E., & Cushman, F. (2015). Philosophers' biased judgments persist despite training, expertise and reflection. *Cognition*, *141*, 127–137. Retrieved from https://doi.org/10.1016/j.cognition.2015.04.015
Singer, P. (1972). Moral experts. *Analysis*, *32*(4), 115–117. Retrieved from https://doi.org/10.1093/analys/32.4.115
Singer, P. (1988). Ethical experts in a democracy. In D. M. Rosenthal & F. Shehadi (Eds.), *Applied Ethics and Ethical Theory, chapter 9* (pp. 149–161). University of Utah Press.
Steinkamp, N. L., Gordijn, B., & ten Have, H. A. M. J. (2008). Debating ethical expertise. *Kennedy Institute of Ethics Journal*, *18*(2), 173–192. Retrieved from https://doi.org/10.1353/ken.0.0010
Táíwò, O. O., & Cibralic, B. (2020). The case for climate reparations. *Foreign Policy*. Retrieved from https://foreignpolicy.com/2020/10/10/case-for-climatereparations-crisis-migration-refugees-inequality/
Weinstein, J. R. (2014). What does public philosophy do? (Hint: It does not make better citizens). *Essays in Philosophy*, *15*(1), 33–57. Retrieved from https://doi.org/10.7710/1526-0569.1488
Wolff, J. (2020). *Ethics and public policy: A philosophical inquiry* (2nd ed.). Routledge

Yancy, G. (2015). Dear White America. *New York Times*. , Retrieved from https://nyti.ms/2ifnyBh
Yoder, S. D. (1998). The nature of ethical expertise. *The Hastings Center Report*, *28*(6), 11–19 Retrieved from http://www.jstor.org/stable/3528262
Zoloth-Dorfman, L., & Rubin, S. B. (1997). Navigators and captains: Expertise in clinical ethics consultation. *Theoretical Medicine*, *18*, 421–432. Retrieved from https://doi.org/10.1023/A:1005884231731