

Karen Jones/François Schroeter

Moral Expertise

Abstract: This paper surveys recent work on moral expertise. Much of that work defends an asymmetry thesis according to which the cognitive deference to expertise that characterizes other areas of inquiry is out of place in morality. There are two reasons why you might think asymmetry holds. The problem might lie in the existence of expertise or in deferring to it. We argue that both types of arguments for asymmetry fail. They appear to be stronger than they are because of their focus on moral expertise regarding all-in judgments about rightness. We reject this emphasis on all-in judgment in favor of an account of moral expertise as typically multi-stranded and domain limited. This account of moral expertise is better able to address the problem of how to identify those who have expertise. It also offers a more nuanced picture of the contrast between accepting a moral claim on one's own and accepting it on testimony.

1. Introduction

We recognize the possibility of expertise in many domains, from theoretical physics to art history, from medicine to plumbing. In each of these areas, there are social mechanisms for identifying experts who then assume a privileged role in our shared epistemic practices. We are willing to call on them when in doubt and to take their opinions to be authoritative or at least to have significantly more weight than the opinions of the non-expert. Morality, however, is thought to be different: here, some argue, there can be no expertise, and if there is no expertise, there can be no experts. Others think that whether or not there is expertise, there is no parallel role in ethics for experts. We should not defer in our moral judgments, and each must reach a decision about what is right according to his or her own best lights.

We begin by offering two contrasting models of what moral expertise would be like, if it exists. According to the first model, moral expertise is a matter of knowing, and knowing how to apply, correct moral principles. The second model, in contrast, takes moral expertise to be a matter of experienced practical judgment. Both models identify expertise with the generalized capacity to form all-in judgments about rightness. We reject this emphasis on all-in judgment in favor of a multi-stranded, domain limited conception of ethical expertise. We then ask: if there is moral expertise, what follows for moral epistemology? This

question is the focus of much recent discussion. What is at stake is the extent to which moral epistemology is social.

2. Two Models of Moral Expertise

It is useful, as a starting point, to consider two diametrically opposed models of what moral expertise, if it exists, would consist in. Which model you will prefer depends on the stance you take about the role of rules and principles in correct moral judgment. According to the first model, moral expertise is a matter of “subsuming the ‘facts’ relevant to a particular problem under a moral theory” (Caplan 1992, 33). In the best case, those with moral expertise will be able to arrange moral principles in a hierarchy and so be able to deduce what ought to be done in any given situation simply by applying the rules. According to this model, central capacities grounding expertise are: the ability to reason well, including critical reasoning skills such as the capacity to identify fallacies; knowledge of the appropriate moral principles, and if there is more than one, of their ordering (Singer/Wells 1984, 200). This knowledge is paradigmatically the kind of thing acquired through the philosophical study of ethical theory; hence this model of moral expertise supports the view that philosophers are more likely to have moral expertise than are non-philosophers.

In contrast with this intellectualist model, the second model, the practical wisdom model, emphasizes perception. It is the model that lies behind contemporary particularist approaches to virtue ethics, such as those advocated by Jonathan Dancy (1994; 2003) and John McDowell (1978). According to this model, the core capacity grounding moral expertise is a capacity to discern the moral salience of considerations in particular contexts. This model is strongly associated with the doctrine of the unity of the virtues. Having a virtue is a matter of possessing the capacity to respond appropriately in deliberation, feeling, and action to the practical significance of that virtue’s signature reason. The honest, for example, grasp the practical significance of something’s being the truth. Suppose that, in addition, we accept the claim that virtues are states that issue in only right action. Given that situations can present more than one of these signature reasons—for example, kindness can appear to conflict with honesty—then we have to accept that you cannot have one virtue without having them all (McDowell 1978, 331–333). Moral expertise is thus a matter of having the virtues, and having the virtues requires having practical wisdom. The deliverances of practical wisdom cannot be captured in a set of rules (McDowell 1978, 338–337; Dancy 1994, 66–71). Practical wisdom is acquired only through habituation in the virtues and experience. On this model, philosophers are no more likely to have moral expertise than non-philosophers. We should look for expertise among the virtuous.

As we have presented them, these two models represent opposite poles of what can be seen as a spectrum of views about the nature of moral expertise; as such, they are caricatures and there is much in between territory. For example, particularists can allow that there is a significant role for intellectual reflection

in criticizing the principles and presuppositions governing our moral thinking (McDowell 1994, 81). At the other pole, any plausible view of moral principles must allow that they do not apply themselves and that it takes judgment, refined by experience, to know how to interpret a situation so that principles can be correctly applied to it.

Both models share the assumption that expertise is primarily about the capacity reliably to deliver action-guiding verdicts that will settle what we are to do in particular circumstances. This focus on action-guidingness emphasizes the integrative role of expertise. If principles can conflict, then expertise must consist not only in the ability to see how they are implicated in particular contexts, but also in the ability to rank these principles so as to get an all-in verdict as to what to do. This explains the push, in the first model, to understanding expertise as requiring knowledge of a unified moral theory. Similarly, this focus on action-guidingness drives the second model's commitment to the unity of virtues. However, both models could drop this assumption if they gave up the idea of confining expertise to action-guiding judgment. After reviewing some objections to the very idea of moral expertise, we will return to examine this assumption, for it turns out that a plausible response to the problem of how to identify who has expertise abandons the thought that expertise is shown only or primarily in all-in judgments about rightness and instead looks for more piece-meal context dependent expertise with particular thick moral concepts, such as 'honest' and 'respectful', rather than with thin concepts such as 'right' and 'ought'.¹

3. Objections to Moral Expertise

The first source of objection to moral expertise is metaphysical. If there are no moral facts, or moral facts simply reflect subjective dispositions, there can be no moral knowledge. If there is no moral knowledge, then there can be no moral expertise. There are a number of well-known arguments purporting to show that there can be no objective moral facts, including the argument from queerness, the argument from explanatory redundancy, and the argument from disagreement. The argument from queerness claims that objective moral facts would have to be completely different from ordinary facts, since they seem to have action-guidingness built into them. Being *sui generis* facts, we would seem to need a distinct epistemic pathway to come to know them (Mackie 1977, 38–42). The second argument claims that there is no need to posit moral facts to explain what Harman calls our moral observations, that is, “our immediate judgment made in response to the situation without any conscious reasoning having taken place” (1977, 6). We can replace explanations that appeal to moral facts by explanations that refer exclusively to non-moral facts and facts about the observer's moral sensibility without loss of explanatory power. Things are

¹ For a famous argument that moral knowledge is possible only with respect to the so-called 'thick' ethical concepts, such as the virtue terms, which contain significant descriptive content, and not with respect to 'thin' concepts, such as right and good, that seem to lack such descriptive content, see Bernard Williams 1985.

alleged to be otherwise in science, where the best explanation of our scientific observations and beliefs requires that we posit unobservable entities.

The third line of argument begins with the observation that there is persistent disagreement about moral matters and claims that the best explanation for this disagreement is that there are no objective moral facts (Mackie 1977, 36–38). Disagreement does not always rule out there being a fact of the matter about which some people are right and others wrong; for example, there is considerable disagreement in most sciences. However, this disagreement can be resolved by rational procedures and convergence can be expected in the long run. Things seem otherwise in ethics, where disagreement appears rationally unresolvable and “seems to reflect people’s adherence to and participation in different ways of life” (Mackie 1977, 36).

None of these arguments are decisive against a realist position, which posits objective moral properties and facts; moreover, there are well known rejoinders (Boyd 1988; Sturgeon 1988). However, even if those rejoinders fail, and we came to agree that there are no objective moral facts, that would still not settle the question of expertise. Only two main metaethical positions are clearly incompatible with the existence of moral expertise, one cognitivist, and one non-cognitivist. Simple subjectivism, a cognitivist position, claims that moral wrongness is reducible to facts about people’s dispositions to disapproval: what makes my utterance, ‘abortion is wrong’, true is that I am disposed to disapprove of abortion. Simple expressivism, a non-cognitivist position, holds that moral judgments, despite their surface assertoric form, are not truth-assessible. Instead, they function to express pro and con-attitudes, like approval and disapproval (Ayer 1936; Stevenson 1944). If moral judgment is simply a matter of reporting or expressing a subjective stance, there is no stable moral subject matter about which some could be mistaken and others correct; hence, there is no moral expertise. Neither of these views, however, are serious contemporary contenders, because they have difficulties adequately explaining the central intuition that judging something right does not make it so. Contemporary expressivists, such as Simon Blackburn (1996; 1998) and Allan Gibbard (1990; 2003), have developed sophisticated strategies for vindicating this intuition within an expressivist framework. These strategies allow them to talk of moral truth in a minimalist, non-correspondence way, and hence to talk of moral knowledge, thus opening space for moral expertise.²

Another line of objection to moral expertise stems from concern about the social role of moral experts. There is a tendency in the literature to slide from talk of expertise to talk of experts (McGrath 2008, 112). Some think of expertise as “exclusionary and restricted” (Archard 2011, 120), thereby assuming that if there is expertise, few people can have it. If many have a capacity then it counts as competence rather than expertise. On this picture, expertise is closely tied to the notion of an expert; but to be an expert is not merely a matter of hav-

² Once a minimalist notion of truth has been established, it is possible to define epistemic norms for acceptance of moral claims as true and norms for evaluating informants as better or worse at judging moral matters. Hence talk of moral knowledge and moral expertise can be accommodated within an expressivist framework (Gibbard 2003, 221–236).

ing expertise. Experts occupy distinct roles in our economy of credibility: their opinion is sometimes to be taken as definitive and always to be taken more seriously than that of a non-expert. 'Expert' is thus not only an epistemic status, it is also a social role. In the moral case, some have objected that this social role is problematic (Caplan 1992; Jones 1999). The problem with a socially sanctioned role of moral expert is political rather than epistemological. In pluralistic democratic societies in which there is significant disagreement in conceptions of the good there will be disagreement over who has moral expertise and so who is fit to serve in the role of moral expert. Moreover, unless experts have a merely advisory role, they risk subverting the democratic process. In a democracy, it matters that everyone be able to accept a policy outcome as legitimate, even if they disagree with it. That outcomes be the result of democratic processes helps secure this legitimacy but is in tension with the idea of giving moral experts more than an advisory role.

Of course, once there is general agreement about the principles or values that are to govern decision-making in a given domain, there is room for recognizing experts in practical matters. The paradigmatic example of this is the law: judges, lawyers, constitutional experts, and so on, can often reasonably claim to be experts regarding how to interpret legal principles. Similarly, experienced members of ethics committees might reasonably claim to have expertise in interpreting how those agreed legal and moral values that provide the parameters for their deliberation apply in a particular context. However, this is not the same as according them the role of genuine moral expert, where the very principles that are to guide decision are themselves up for grabs.

The political objections to a socially sanctioned role of moral experts are not, however, objections to moral expertise in itself. They also rest on an account of expertise that there is good reason to reject. While experts must be few and far between, given their social role, expertise can be widely shared. To borrow a list from Arthur Caplan, commonplace expertise includes "how to drive a car, mail a letter, exhibit good manners, use a post office, play games" (Caplan 1992, 30–31). Similarly, we might suppose that a great deal of moral expertise is commonplace among ordinary adults. The issue of expertise and of the socially sanctioned role of experts must be separated.

The traditional version of the argument from disagreement rejects moral expertise because, denying there are moral facts, it rejects moral knowledge. A new epistemic argument moves directly from disagreement to moral skepticism, or at least to skepticism about those controversial moral questions we might want to call on expertise to help us answer (McGrath 2008). The argument rests on the problem of how to identify experts. However, unlike the previous argument, which confused expertise and the social role of experts, this argument addresses the first-person concern of how to rank oneself with respect to expertise. If this cannot be done, self-trust is undermined, and with it one's claim to knowledge.

McGrath begins from the *prima facie* plausible principle that one cannot know that *p* if *p* is denied by someone of whom "you have no more reason to believe he or she is in error than you are" (91). McGrath then argues that all 'hotly' contested moral questions, such as the moral permissibility of the

death penalty, abortion, and meat eating, are of this form. People who are our epistemic peers—that is, people whose judgment we have no reason to suppose is more likely to be in error than ours—stand on both sides of these issues. Even if there are moral facts that determine the answers to these questions, and even if one side but not the other is reliably connected to these facts, the existence of peer disagreement undercuts either side’s claim to knowledge (92). McGrath argues that the ways in which we identify those with expertise in other domains do not apply in the moral case. In some fields we can check expertise independently; for example, we can check the expertise of weather forecasters by seeing if their predictions hold, of engineers by seeing if their bridges hold. In addition, for a wide class of experts, from plumbers to physicists, there is agreement about appropriate training. But, claims McGrath, neither of these apply in the moral case.

The problem McGrath identifies is important; however, it is not clear that it supports an extensive moral skepticism, nor that it is as undermining of the notion of expertise as she seems to think. It is perhaps not surprising that we cannot claim to know the answer to many currently contested moral questions, opposing sides of which are defended by well-informed people of intelligence and integrity. Moreover, lack of knowledge here is compatible with knowledge about a great many moral claims, including some claims that are not part of a general moral consensus in the way a claim like ‘slavery is wrong’ is. These include claims about the moral status of women and people of color, and, potentially, any claim for which we can come up with a credible theory explaining sources of error. McGrath seems to think that there is a danger many explanations that purport to show why one side is more likely to be in error than the other will be question-begging or circular (99). It is true that all such explanations will be theory-dependent and so will presuppose further contestable claims in their explanations. But this is a general feature of theory acceptance, including theory acceptance in science, and is not the kind of special problem for ethics that could support asymmetrical skepticism (Boyd 1988, 189–191; Jones 2005, 76–7).

When it comes to disagreement about all-in verdicts such as whether an action is right or wrong—that is, when it comes to correctly applying our thin moral concepts—it is harder to assess the comparative epistemic qualifications of disputing parties than it is when it comes to claims using thick moral concepts such as those expressed by the virtue terms or terms for specific moral values such as respect. None of this is to say that there cannot be moral expertise about all-in moral judgment or that identifying those with such expertise is impossible; however, we currently do not have the kind of story about how such generalized expertise is acquired that we have regarding the acquisition of more particular types of expertise. McGrath only considers all-in judgments, and yet they form just a fraction of our moral judgments. We not only form verdicts about what is right and wrong, but also about what is just and courageous, about what kinds of moral considerations are relevant to settling an issue, about the nature and significance of different values, such as the nature and significance of respect, and so on. The realm of moral judgment is complicated and various.

It is easier to understand what kinds of experience and training would put an agent in a better epistemic position with respect to the nature and significance of these specific values or virtues. Think, for example, of the moral understanding of the need and frailty of human beings that can emerge from a life, such as Mother Teresa's, spent caring for the dying destitute. Such lives might not present equivalent lessons in the value of social justice as lives spent resisting political repression (Jones 1999, 65). Sometimes moral expertise is patchwork rather than generalized. Moreover, the same value can manifest itself differently and interact differently in different contexts; for example, justice manifests itself differently in the family than in the political sphere. These different domains present their own characteristic patterns in the combinations of moral considerations that tend to occur within them. Different agents can have greater levels of experience with the ways a value works in one context than with the ways it works in another. This opens up the possibility of domain specific expertise regarding all-in judgment.

Either model of expertise can accommodate fragmented, multi-stranded, domain limited moral expertise, once it loosens its focus on integrative judgment. The different models will put the point somewhat differently: on the theory-based model, some moral agents might have expertise about the ways in which principles apply in particular contexts, others about how they rank, or rank in a given domain. On the virtue-based model, some agents will be further along the path towards achieving some virtues than others, some will have a more refined capacity for detecting the salience of ethical considerations in some domains than in others. Given that the world of value is complex and the capacities needed to navigate in it many and various, it is more realistic to expect that such human moral expertise as exists would take patchwork form rather than the idealized form hypothesized in the simple statement of the two models. Recognizing the plurality of forms of expertise has implications for the epistemic significance of expertise, which is our topic in the next section.

4. Epistemic Implications of Moral Expertise

Let's assume there is moral expertise—whether patchwork, limited, and domain-specific, or broader. What does this mean for moral epistemology? There is a minimal sense in which recognizing moral expertise pushes us towards a social epistemology. When we deliberate about practical matters, we should deliberate together, and this is not just a matter of two heads being better than one, but rather a matter of making sure that we call on all relevant expertise in resolving practical matters. Our moral epistemology is thus already more social than it would be on an account which supposes all normal adults have the relevant moral expertise and capacities and must rely on each other only to check that they are reasoning well and not overlooking anything. But does the existence of moral expertise support a richer social epistemology, of the kind we find in science? A richer social epistemology includes epistemic divisions of labor, which in turn support assigning differential epistemic positions within a community of

inquiry, positions that bring with them different rankings on the credibility scale, relations of deference, and roles for expert testimony and advice. Much recent work has been devoted to these questions. On the one hand, there are those who think that there is no deep difference in well-structured moral epistemic communities and other communities of inquiry, including science (Coady 1992; Jones 1999). On the other hand, there are those who think that the moral case is special (Nickel 2001; Hopkins 2007; Hills 2009; McGrath 2008; 2011). In the remainder of this paper, we will investigate the major arguments in favor of an asymmetry thesis, focusing on moral testimony. These divide into two kinds: epistemic arguments that purport to show that moral knowledge cannot be transmitted via testimony, though knowledge about ordinary facts can, and practical arguments that purport to show that while moral knowledge might be transferable, we should not get our knowledge in this way. The second, more interesting, line of argument will be our main target.

There are two main epistemic arguments against moral testimony, even granting the existence of moral expertise: the practicality and the significance argument. According to the practicality argument, the close link between moral judgment and action makes moral knowledge a matter of knowledge how, rather than knowledge that. Testimony can transfer knowledge that, but not knowledge how; hence it is useless in the moral case (Anscombe 1981, 47–8; Hopkins 2007). This argument overstates the role of knowledge how in morality. Though it is plausible that our moral competences include some knowledge how, it is implausible to suppose that they are not in large part constituted by propositionally articulated beliefs that could be transferred by linguistic means and so could be acquired from testimony (Hopkins 2007, 618–20).

The significance argument begins from the observation that morality matters: what is at stake includes benefits and harms, life and death decisions, and whole ways of living lives and organizing societies. If morality is so very important, then it would seem we have a responsibility to be especially careful to get it right, and this in turn might suggest a responsibility to think things through on our own, rather than relying on the testimony of others. However, this objection is weak: beliefs about non-moral matters, such as cancer treatments, can be every bit as important as beliefs about moral matters. Moreover, the more important things are, the less confident we should be that our capacities are up to the job. Far from supporting an individualist approach, the importance of moral matters suggests we should draw on all available resources for resolving them, including any available resources of expertise (Jones 1999, 57, 66–7).

If the problem with moral testimony is not that it cannot transfer moral knowledge, then perhaps it is that we have other strong reasons not to get our moral knowledge via testimonial means. The deficiency in moral beliefs acquired through testimony might not be epistemic, but rather related to deficiencies in the moral worth of actions performed on its basis, or deficiencies in the moral worth of agents who need to rely on it. The most systematic elaboration of this line of argument is Hills 2009.³

³ See also Nickel 2001, for an earlier version of the argument that actions grounded in

Hills argues that agents who rely on moral testimony lack moral understanding. Because they lack this understanding their actions and they themselves as agents are morally deficient. Most discussions of testimony focus on the question whether one can come to know that *p* through testimonial transmission. But knowing that *p* is different from knowing why *p*, and knowing why *p* is different again from *understanding* why *p* (Hills 2009, 100–106). To understand why an action is morally wrong, you must understand the reasons why it is wrong. Understanding requires more than knowing that there is a reason why it's wrong, or even knowing what that reason is, you must also grasp the support relation linking the consideration that is the reason to the action for which it is a reason. This support relation cannot be grasped in isolation from understanding the relative importance of that consideration in relation to other considerations that might bear on that action as well as how that consideration might support other actions in different contexts. Thus, this understanding cannot be easily isolated in a well defined set of local knowledge-that claims. Hills argues that understanding why *p* involves the abilities to follow someone's explanation as to why *p*, give an explanation in your own words, draw the conclusion that *p* from the considerations that support it, make related inferences in similar but not identical cases, and give the right explanations in relevantly similar cases, even where those explanations differ somewhat (102–103). Clearly you can know why *p* (that is, know what reason supports *p*) without having these further capacities.

Moral understanding has epistemic value insofar as it equips us better to contribute to shared inquiry and insofar as it helps us get things right in a range of related cases (106–107). But, according to Hills, the central value of moral understanding is itself moral: moral understanding affects the moral worth of agents and of actions. On the standard conception of virtue, a virtuous agent not only reliably does the right thing, she possesses moral understanding; indeed, given the doctrine of the unity of the virtues, she possesses fully general moral understanding. To the extent that one lacks understanding, one lacks virtue.

More controversially, Hills claims that the moral worth of an action depends on the agent possessing an understanding of the reasons that make the action morally right. The reasons for which an action is done make a difference to what is done. That there is a distinction between “‘outward’ actions” and “actions-for-reasons” (118) is familiar from Kant (1991); however, what is distinctive about Hills' development of this thought is the requirement not only to know what the reason is, and to act on its basis, but to act with the kind of orientation towards that reason that presupposes the broader abilities required for understanding. That reasons for acting matter is plain enough: though Kant's self-interested shopkeeper does the right thing, he does not do it for the morally right reason and this deprives his action of moral worth. When someone does the right thing simply because they have been told that it is right by a trustworthy informant, it seems that, likewise, their action lacks moral worth. While they are concerned about doing the right thing, they do not have access to the reasons that make their action right and so cannot act for them. At best their reason is, “it is

testimony lack moral worth. A related idea, deploying the notion of a ‘right’ to hold moral beliefs, is found in Hopkins 2007.

right, he told me". For the sake of argument, let's grant this point to Hills: there is something deficient about the actions of a person who acts from bare testimony about all-in rightness. Their actions are better than those of someone who doesn't care about the moral quality of their actions, but they are not responsive to that which makes the actions right.

Hills' point is not yet established, however, since not all moral testimony is testimony using thin moral concepts such as rightness. Much moral testimony concerns what reasons are and are not relevant in a particular context. As noted above, this kind of testimony is paradigmatic rather than the marginal case as it is here that we are more likely to be able to identify those who have more expertise than we do and so pick good informants. The real test for Hills' view are cases where the testifier transfers their knowledge of the right-making reasons to the recipient of their testimony, as when, for example, someone says that it is right to give money to people living in poverty because justice, not just charity, requires it. In cases of this kind, the recipient of testimony is concerned about doing the right thing and has chosen a knowledgeable informant on the basis of whose testimony she now knows that an action is right, and knows the reason why it is right but still lacks the abilities required to understand why it is right. Hills claims that "more is required for morally worthy action: you need to act for the reasons that make your action right" (Hills 117). The puzzle, however, is why having been granted testimonial access to the reasons that make your action right you are still not acting for those reasons. Hills claims you are not because you are not properly oriented towards those reasons, but that is disputable. You are not independently oriented towards those reasons—you require the assistance of someone else to turn your attention and motivation in their direction—but still it seems, with this assistance, you can come to respond to them. What is at issue here is whether an agent needs the kind of rich orientation towards reasons that requires the understanding characteristic of virtue before she counts as responding to them. Hills has not given an argument that this is necessary. We might think that those with understanding have a stronger counterfactual link to reasons than those who lack it: that is, were the situation to be somewhat changed they would still be capable of responding appropriately to their reasons. However, reliability can be achieved second-hand, provided an agent correctly chooses their informants. Moreover, it is unclear why responding to reasons on one occasion would require the ability reliably to respond to similar reasons in different cases. In sum: acting on the basis of someone else's all-in judgment about what action is right without any explanation of what the reasons are might indeed be incompatible with the moral worth of that action; but for all that has been argued, when knowledge of the right-making reasons is also transferred, and one acts on that basis, the action has moral worth.

It might be thought that this is a small concession, and one Hills can readily make. However, the claim that actions grounded in testimony lack moral worth is central to her argument for drastically limiting the scope of acceptable testimony in matters of morality. Hills claims that it is incumbent on us, as moral agents who must act well and not merely rightly, that we develop and display our understanding. Testimony cannot transfer understanding and so when we rely

on testimony we neither develop nor display our understanding and yet morality requires us to do both. For this reason, she claims that we should rely on testimony only when, due to immaturity or some deficiency in us, understanding is precluded (2008, 123). We are not similarly obligated to restrict reliance on testimony in other domains where the moral worth of our action is not at issue.

This argument fails if actions can have moral worth even when the agent is deficient in understanding. There is no requirement that each action be based on understanding, nor that we take each and every opportunity to develop that understanding. Even if we grant that there is a general requirement to develop our moral understanding, we cannot move from that general requirement to a case-by-case requirement. Furthermore, it is disputable whether there is in fact a general requirement to develop moral understanding, even if we concede, as we should, that full virtue presupposes such understanding. The conception of full virtue as involving the ability to integrate moral considerations into all-in rightness judgments regardless of the domain at issue is highly idealized.⁴ As we have argued, a realistic picture of human agency admits human moral excellence is more likely to be more patchwork and domain limited than this ideal recognizes. The fact that human beings are unable to achieve the ideal of full virtue does not, in itself, count against the worth of that ideal. There can be a point to striving to achieve an ideal even though it is not reachable: by trying we may do better than we would if we set our sights lower.⁵ But sometimes, unachievable ideals are the enemy of achievable second-bests. We can see this in the ordinary epistemological context where an ideal of cognitive autonomy that amounts to epistemic individualism is not only unachievable given the testimony saturated nature of our knowledge, but trying to achieve it takes our attention away from the pressing questions that finite knowers invariably face such as to whom to defer, regarding what. Likewise, we may do better as moral agents, by acknowledging the limitations in our abilities to develop our moral capacities in all domains and calling on the help of others. Hills' proposed drastic limitations on the place of moral testimony close off this possibility.

Once we reject Hills' injunction to limit testimony only to cases where understanding is precluded, the question of how to be a wise recipient of moral testimony becomes central (Jones 1999). What responsibilities must be discharged before our deference counts as responsible? The issue is not just whether or not one chooses to defer to someone who in fact has the relevant expertise: more is at stake than just success. One might defer to someone who can pass on the relevant knowledge, and yet one's deference be both epistemically and morally irresponsible, as, for example it will be, if one has good reason to believe that the other person does not have the relevant knowledge.

⁴ As is the related principle-based conception of the ideal moral agent as able to apply ranked rules or principles in all possible situations.

⁵ The distinction between target and regulative ideals is found in Louise Antony 1995. See also John Doris 1998 for an argument that empirical research shows we would be better off giving up the traditional conception of a virtue even as a regulative ideal.

The question of what it takes to be a wise recipient of testimony is obscured by the choice of examples that dominate the literature on moral expertise. Hills, for instance, gives us the example of Ron (first introduced in Arpaly 2003, 74), who, by hypothesis, cares about doing the right thing, and chooses as his informant a rabbi, who, again by hypothesis, is a trustworthy testifier. On this basis, Ron comes to know that killing Tamara is wrong and that the reason it is wrong is because Tamara is a person. The stipulative nature of the example and its corresponding lack of psychological realism obscures two problems: it sidesteps investigation of how to choose testifiers and it posits a moral psychology that is so bizarre and unrealistic we cannot see how anybody with that psychology would be equipped to discharge the responsibilities of being a wise recipient of testimony. We are at a loss to understand how someone—let alone someone who cares about doing the right thing—could be so clueless about moral reasons that they cannot see that someone’s being a person is a reason not to kill them. In certain contexts, it can be difficult to understand the reason-giving force of personhood—does it, for example, count against or in favor of voluntary euthanasia?—but that personhood bears on the permissibility of murder seems to be within the grasp of anyone capable of possessing the concepts ‘personhood’ and ‘moral rightness’. Moreover, once we think seriously about what it would be like to be so clueless as to require testimony about something this basic, we should rightly wonder about Ron’s ability to pick out a trustworthy informant. Why did Ron choose a rabbi, who by stipulation is a good source of moral knowledge on this matter, rather than an extremist who recognizes value in some lives but not others?

What does it take to be a wise recipient of moral testimony? Just as there are two ways of approaching this question regarding testimony in general, there are two ways here. On the first approach, there is a default assumption in favor of accepting testimony that requires positive evidence to rebut (Hinchman 2005a); while on the second, more demanding approach, one needs positive evidence of the trustworthiness of the testifier before it is responsible to accept testimony (Jones 1999).⁶ In the practical context where important matters are at stake, and where we know moral views are subject to the distorting influences of self-interest and ideology, there is a strong case to be made in favor of the second approach.

The second approach requires a wise recipient of moral testimony to have a sense of their own zones of comparatively greater and lesser moral competence, so that they will know where they might need to call on the help of others and where they can themselves claim expertise. Having this reflective self-understanding amounts to having a rough map of your own competence, which in turn requires a sense of the kinds of things that enable and disable moral knowledge.⁷ Not only must you have this map, you must apply it in choosing when to defer and

⁶ Philosophers defending a general default permission to accept testimony include Burge 1993, Coady 1992 and Reid 1983. For a more cautious approach, see Fricker 1994.

⁷ Reasonable reflective self-understandings are unlikely to be able to be arrived at on one’s own. We use the judgments of others to calibrate our own epistemic methods and we form and test theories about what makes for reliability in an area together.

who to defer to. The capacities for wise self-trust and distrust are thus not distinct from the capacities needed to be wise in your trust of others.

The picture of a wise recipient of moral testimony that emerges is thus incompatible with the kind of ‘blank slate’ that an agent like Ron appears to have. The more you are able to understand at least some areas of moral value, the better you will be at picking out who to trust. Realistic examples of the wise receipt of moral testimony occur against a background of at least some moral expertise on the part of the recipient, which provides understanding of at least some reason-relations against which the other person’s proposal can be seen to make sense. What is taken to be part of this background can depend on context and can change as moral understandings change over time. This background simultaneously places rough constraints on the content of testimonial claims that it would be responsible for us to accept and places significant constraints on whom we can reasonably see as potential sources of moral knowledge. For example, if someone were to tell us that the fact that more people will count blades of grass is a morally significant reason for preferring one course of action over another (Rawls 1970, 432), they would have said something so incredible that they would thereby have shown themselves to be untrustworthy as a source of moral knowledge. In the absence of any points of contact between the moral sensibility of testifier and recipient, a testimonial transaction would be bound to fail the responsibility test.

The requirement that a wise recipient of moral testimony assess the trustworthiness of another in moral matters does not mean that the wise must be able independently to assess the content of the claim accepted on testimonial grounds. The primary object of assessment is the *person*, not the claim. Moreover, there is no requirement that the assessment of the person be conducted solo—often we will draw on the knowledge of others to help us identify those who might have the relevant kind of expertise. It is tempting to see a stark contrast between testimonial deference and co-deliberation in which, after critical assessment, we take other people’s judgment into account in our deliberation as something to be accepted on its merits. But once we bring the rich background against which wise testimonial transfer occurs, this sharp contrast fades. As we discharge our responsibilities as recipients of moral testimony, we calibrate others as potential sources of moral knowledge and locate them on our map of our own and other people’s moral competencies and comparative expertise. There are not two options here, full-blindness versus complete independent assessment of the merits the case. Even if we reject full-blindness, as we should, we still retain a robustly social epistemology with room for epistemic divisions of labor and relations of deference.

Bibliography

- Antony, L. (1995), Sisters, Please, I’d Rather Do It Myself: A Defense of Individualism in Feminist Epistemology, in: *Philosophical Topics* 23, 59–94
- Anscombe, E. (1981), Authority and Morals, in: *The Collected Philosophical Papers of G. E. M. Anscombe, Vol. III, Ethics, Religions and Politics*, Minneapolis, 43–50

- Archard, D. (2011), Why Moral Philosophers Are Not and Should Not Be Moral Experts, in: *Bioethics* 23, 119–127
- Arpaly, N. (2003), *Unprincipled Virtue*, Oxford
- Ayer, A. J. (1946), *Language, Truth, and Logic*, New York
- Blackburn, S. (1996), Securing the Nots: Moral Epistemology for the Quasi-Realist, in: Sinnott-Armstrong, W./M. Timmons (eds.), *Moral Knowledge? New Readings in Moral Epistemology*, Oxford, 82–100
- (1998), *Ruling the Passions: A Theory of Practical Reasoning*, Oxford
- Burge, T. (1993), Content Preservation. in: *The Philosophical Review* 102, 457–488
- Boyd, R. (1988), How to Be a Moral Realist?, in: Sayre-McCord, G. (ed.), *Essays in Moral Realism*, Ithaca, 181–228
- Caplan, A. L. (1992), *If I Were a Rich Man Could I Buy a Pancreas? And Other Essays on the Ethics of Health Care*, Bloomington–Indianapolis
- Coady, C. A. J. (1992), *Testimony: A Philosophical Study*, Oxford
- Dancy, J. (1993), *Moral Reasons*, Oxford
- (2004), *Ethics without Principles*, Oxford
- Doris, J. (1998), Persons, Situations, and Virtue Ethics, in: *Noûs* 32, 504–530
- Fricker, E. (1994), Against Gullibility, in: Matilal, B. K./A. Chakrabarti (eds.), *Knowing from Words*, Boston, 125–161
- Gibbard, A. (2003), *Thinking How to Live*, Cambridge/MA
- (1990), *Wise Choices, Apt Feelings: A Theory of Normative Judgment*, Cambridge/MA
- Harman, G. (1977), *The Nature of Morality*, Oxford
- Hills, A. (2009), Moral Testimony and Moral Epistemology, in: *Ethics* 120, 94–127
- Hinchman, E. (2005a), Telling as Inviting to Trust, in: *Philosophy and Phenomenological Research* 70, 562–587
- (2005b), Advising as Inviting to Trust, in: *Canadian Journal of Philosophy* 35, 355–386
- Hopkins, R. (2007), What Is Wrong with Moral Testimony?, in: *Philosophy and Phenomenological Research* 74, 611–34
- Jones, K. (1999), Second-hand Moral Knowledge, in: *Journal of Philosophy* 96, 55–78
- (2005), Moral Epistemology, in: Jackson, F./M. Smith (eds.), *The Oxford Handbook of Contemporary Philosophy*, Oxford, 63–85
- Kant, I. (1991), *Groundwork of a Metaphysics of Morals*, trans. H. J. Paton, London
- Mackie, J. (1977), *Ethics: Inventing Right and Wrong*, Harmondsworth
- McDowell, J. (1979), Virtue and Reason, in: *The Monist* 62, 331–350
- (1994), *Mind and World*, Cambridge/MA
- McGrath, S. (2008), Moral Disagreement and Moral Expertise, in: Shafer-Landau, R. (ed.), *Oxford Studies in Metaethics* 3, 87–107
- (2011), Skepticism about Moral Expertise as a Puzzle for Moral Realism, in: *Journal of Philosophy* 108, 111–137
- Nickel, P. (2001), Moral Testimony and its Authority, in: *Ethical Theory and Moral Practice* 4, 253–266
- Rawls, J. (1971), *A Theory of Justice*, Cambridge/MA
- Reid, T. (1983), *Inquiry and Essays*, Indianapolis
- Singer, P./D. Wells (1984), *The Reproductive Revolution*, Oxford
- Stevenson, C. L. (1963), *Facts and Values: Studies in Ethical Analysis*, New Haven
- Sturgeon, N. (1988), Moral Explanations, in: Sayre-McCord, G. (ed.), *Essays on Moral Realism*, Ithaca, 229–255
- Williams, B. (1985), *Ethics and the Limits of Philosophy*, London