

# MORAL EXPERTISE: JUDGMENT, PRACTICE, AND ANALYSIS\*

BY JULIA DRIVER

There is some debate over whether or not moral experts exist. The debate on whether or not moral experts exist has tended to conflate at least three different ways in which someone might be thought to be a moral expert. My primary thesis in this essay is that there are moral experts. Further, there are at least three distinct forms moral expertise can take: there is the expert judge, who does a better job of arriving at true moral judgments, the expert practitioner, who acts morally well more than others, and the expert in moral analysis who has greater than normal insight into the nature of morality (in some respect). In supporting these claims, I will make use of an analogy between language and morality. Such an analogy has been employed in discussing other features of morality, and it has a good deal of intuitive appeal. My view is that it can also shed light on the (at least) three ways in which one can be a moral expert.

Before addressing this division, however, I would like to consider the issue of why there is some initial skepticism regarding the existence *at all* of moral experts.

## I. WHY THE SKEPTICISM ABOUT MORAL EXPERTISE?

If there are moral experts, they would be persons who are better able to make or are more reliable at making correct moral judgments relative to some other group. Many writers on moral expertise have cited Gilbert Ryle as the primary source of skepticism regarding such expertise.<sup>1</sup> Ryle famously noted that there were no moral experts, because there is nothing for them to be expert *about*. He utilizes a variety of arguments, but the most famous is apt to be the one in which he notes that we do not (and, it seems, cannot) forget the difference between right and wrong once learned, though we do (and can) forget other things that we learn, such as the capitol of France and the year Einstein won the Nobel Prize. For

\* An earlier version of this paper was presented at a Workshop on Moral Expertise presented at Princeton University, April, 2011, organized by Sarah McGrath and David Estlund. I thank the organizers and other participants at the workshop as well as the other contributors to this volume for their very helpful feedback on the paper, and particularly Jamie Dreier for his excellent written comments.

<sup>1</sup> Gilbert Ryle, "On Forgetting the Difference Between Right and Wrong," in *Essays in Moral Philosophy*, ed. A. I. Melden (Seattle, WA: University of Washington Press, 1957).

Ryle, this indicated that morality is not a matter of knowing, but of feeling. There is nothing cognitive about it. Since it is not a matter of knowing, then no expertise is possible. It is not possible for one to know more since *no one* knows anything at all with respect to morality. Eric Wiland has pointed out that Ryle's argument is much too vague. While it is true that it would be very odd indeed for someone to forget the difference between "right" and "wrong" it would also be odd to forget the difference between "true" and "false," and yet that surely does not show that the difference between "true" and "false" is one of *feeling*.<sup>2</sup> To avoid this criticism Ryle would need to make the charge more precise, such as, one cannot forget that (1) "Torturing kittens for fun is wrong" once one has learned it. This does not seem nearly as clear-cut. The lack of clarity here, I believe, depends upon competing views regarding what is involved in really *learning* something about morality: so, in order to learn (1), do I need to simply *accept* it as a true claim, or must I also *appreciate* its truth? Those who view moral knowledge as something that cannot be transferred *simply* via testimony of a reliable source, tend to opt for the latter view. They hold something like the following claim on nontransmission:

(NT) Moral knowledge cannot be transferred *simply* via testimony.

It is important to get the claim right. Proponents of (NT) do not claim that one cannot learn about morality from others. For example, Jessica may come to the conclusion that eating meat is wrong on the basis of learning not only about how animals are treated in factory farms, but also on being made vividly aware of certain moral truths in an ethics class, for example. But this would not be a counterexample to (NT) as long as Jessica comes to appreciate *why* eating meat is wrong, not just *that* it is wrong. The sort of cases that really plumb the intuitions behind (NT) are those that involve someone coming to a conclusion, for example, that torture is wrong simply on the basis of someone he trusts telling him so. For full knowledge, they claim, that is just not enough. It is not enough because that person, unlike Jessica, just does not "get it." One way to unpack the intuition is to hold that moral knowledge is not simply correctly accepting a moral truth on the basis of nonfoundational reasons (such as the testimony of a reliable source), nor on the basis of simply accepting a reason as a good one, it *also* involves feeling a certain way. This does not mean that it is totally a matter of feeling, as Ryle would seem to have it. It simply means that for full moral knowledge of a proposition, the person must appreciate its truth, not simply accept it as true even if one has very good reasons for accepting it as true. There are basically two ways this can be understood. The first involves holding that in order to properly have a

<sup>2</sup> Eric Wiland made this point during the Workshop on Moral Expertise held at Princeton University, April 2011.

moral belief at all one needs to appreciate it as true (even if it turns out not to be). This view is implicitly committed to a form of internalism: moral belief is motivating, and motivation is affective; thus there must be an affective quality, a felt quality of appreciation, let's say, in holding the moral belief. The worry is that when one acquires a "belief" via pure testimony that felt quality is missing, so it does not qualify as an actual belief. The belief is qualified in some way that makes it fall short of full commitment to the *norm* expressed in the testimony. A person might view a psychopath's acceptance of "Killing is morally wrong" as a moral quasi-belief; it is true that when asked about killing, the psychopath would reliably be able to say that it is morally wrong. There is certainly a sense in which the psychopath believes that it is morally wrong. But lacking the felt quality of appreciation, the moral aspect seems to be missing. This is one way to understand the worry. On this view, too, moral testimony differs dramatically from testimony regarding purely descriptive facts. A belief about the height of the pyramids, all by itself, is not motivating.

Another way to unpack the intuition is to grant that one could have a moral belief that is not motivating (and thus grant the externalist possibility), but then hold that this belief cannot rise to knowledge without the appreciation of the reason that underlies justification of the belief. This treats *appreciation* of an underlying moral reason as necessary for the belief to be justified. But this itself is wrong since it conflates two distinct issues. There is having a justified belief that the proposition that expresses a moral claim is true. Then there is *appreciating* it. But note that this debate centers on what counts as sufficient justification for moral knowledge. One can grant that there is some affective ingredient to the mix without losing the cognitive aspect, and this is what is sufficient for the possibility of expertise. Peter Singer, responding to Ryle, noted the ways in which moral philosophers might possess this expertise relative to members of the population as a whole.<sup>3</sup> For example, their training leads them to make finer discriminations, so they have a better understanding of what, morally, is at stake in a given situation. What they are gaining is increased knowledge, not simply finer emotional responses. Thus, the skepticism is unwarranted. Whatever one's views on the necessity of affective justification for moral belief, moral belief is cognitive. There are moral experts, though it may be quite difficult to identify them and justify trust in them.<sup>4</sup>

Alison Hills unpacks the intuition by noting the distinction between moral knowledge and moral understanding. The person who accepts that torture is wrong simply on the basis of testimony has knowledge (given the reliability of the source), but he will not have understanding simply on the basis of that testimony.<sup>5</sup> On this view, then, again, there are experts

<sup>3</sup> Peter Singer, "Moral Experts," *Analysis* 32, no. 6 (1972): 115–17.

<sup>4</sup> I discuss this issue of trust more fully in "Autonomy and the Asymmetry Problem for Moral Expertise," *Philosophical Studies* 128 (2006): 619–44.

<sup>5</sup> Alison Hills, "Moral Testimony and Moral Epistemology," *Ethics* 120 (2009): 94–127.

and those experts can transmit moral knowledge; and the contrary intuition is defused by noting that the person who truly only relies on testimony is still lacking something morally significant. So, given a rejoinder to skepticism both about moral expertise and testimony, what is a moral expert?

## II. MORAL EXPERTISE

Again, on the view that I develop here, I start with the notion that a moral expert in *judging* is someone who can better make correct moral judgments than can others. In the case of the agent, a moral expert is one who acts morally better than others, and in the case of analysis, a moral expert is more skilled at understanding morality than members of the relevant contrast class. It is important to keep these three ways in which someone can be a moral expert distinct, otherwise criticisms might miss the correct target. In an early essay on moral expertise, Béla Szabados discusses who can best give moral advice and is critical of Peter Singer's view that moral philosophers will do a better job of this than others.<sup>6</sup> But different ways in which one can be an expert—as Szabados notes himself—get conflated. Moral philosophers may well have skills that make them better at moral analysis, for example, and spotting the things that make a given act immoral and distinguishing those factors from the morally irrelevant, and so on. But this does not necessarily translate into better moral action. Indeed, one can easily imagine someone who is brilliant in her understanding of ethical theory, and yet has difficulty interacting well with other people.

Keeping the above distinctions in mind, there are two additional issues that need to be considered independently. The first is what it is to be an *expert* (as opposed to a naïf), the second is what it is to be a *moral* expert (as opposed to some other type of expert).

On the first issue, the view that I find most plausible is contrastive in the sense that a claim that someone is an expert is understood relative to a contrast class. There is a sense, on this view, in which all normal adult human beings are moral experts because they are experts relative to the class of five-year-old children. Developing expertise in making moral judgments and engaging in moral action takes some practice and skill development. This is *not* a claim of necessity about expertise. It is possible that there are five-year-old moral savants. But what is typically the case is that human beings develop in *social* situations that require the development of skills to navigate social interactions in a morally appropriate manner. Of course, the most interesting contrast—that is, what people are generally interested in when discussing moral expertise—is not the contrast with the total naïf. Rather, the interesting contrast is that with “nor-

<sup>6</sup> Béla Szabados, “On Moral Expertise,” *Canadian Journal of Philosophy* 8 (1978): 117–29.

mally performing adults.” The question of moral expertise should be understood, then, relative to that contrast. Thus, the question of whether or not John is a moral expert contrasts his judging abilities with those of normal adults: “Is John a moral expert (rather than simply a normally performing judge, agent, analyzer)?”

In the psychology literature, experts tend to be identified as people who meet the following criteria:

1. Experts have devoted at least ten thousand hours to perfecting their decision-making or problem solving abilities *within a particular area*.
2. Experts are capable of more nuanced representations.
3. Because of criterion 2, experts actually perceive the world differently.
4. Experts are better able to sort—that is, they can detect mistakes much better than the naïf
5. Experts have different skill sets: “The expert tries to match the problem with problems held in memory. . . . Unlike novices, experts know what knowledge to access, which procedures to apply, how to apply them, and when it is appropriate.”<sup>7</sup>
6. Expertise is domain specific, so the skill sets will be exemplified in very specific contexts. An expert with a wrench may not be an expert with a power saw.
7. Automaticity: experts tend to make decisions very quickly, with little or no deliberation.

However, some of these conditions are poorly suited to framing a *conceptual analysis* of “expertise”—rather, they point to important developmental conditions, for the most part, that are contingent conditions of expertise in human beings. Consider the first. If ordinary people develop moral expertise it does not seem to work this way since very little of our moral experience as young children involves carefully deliberating about what to do. Children are immersed in a social world in which they are confronted with commonplace moral problems—cooperation, sharing, kindness, and so forth. They need to be able to properly navigate these problems and they rely on some instruction, as well as observation, to figure out what they are supposed to do. The norms of interaction are inculcated this way. However, in the case of expertise in *analysis* the commitment of ten thousand hours seems to be a plausible condition.

To spell out the issues more clearly, consider an analogy with the ways someone could be a linguistic expert. Again, analogies between morality and language are independently appealing. Pursuing the language anal-

<sup>7</sup> Darcia Narvaez and Daniel K. Lapsley, “The Psychological Foundations of Everyday Morality and Moral Expertise,” in *Character Psychology and Character Education*, ed. D. Lapsley and C. Power, (Notre Dame, IN: University of Notre Dame Press, 2005), 140–65.

ogy, we may hypothesize that there is some natural disposition on the part of human beings to pick up on these norms and that these dispositions incline us to appropriate action.<sup>8</sup> In the absence of such dispositions we may hold that the development of expertise does require a good deal of deliberative guidance, but that such guidance is not necessary in the development of basic expertise in language and moral judgment. Again, using the language development analogy, adult human beings are linguistic usage experts relative to small children. Language, like morality, is norm governed. In using language properly, one displays the ability to conform linguistic behavior to those norms. Linguistic experts in the sense of those who speak and/or write well, use language better than others. But it is also true that a normal human adult will be a linguistic expert relative to the class of children by virtue of the fact that a normal human adult speaks significantly better than a young child. Further, a normal human adult can identify mistakes in usage better than young children. This is another hallmark of expertise. An expert can *sort* better and worse examples more accurately than the comparative non-expert.

There is evidence that human beings do possess a kind of moral grammar that might underlie moral dispositions, as linguistic grammar underlies linguistic abilities.<sup>9</sup> John Mikhail, for example, discusses a hypothesized Universal Moral Grammar (UMG) that can account for seemingly universal features of moral cognition that are reflected in moral judgments, such as judgments that murder is wrong and judgments that causation is necessary for (at least legal) responsibility.<sup>10</sup> One need not accept all features of UMG to hold that the language/morality analogy is fruitful. But it does support the analogy to show that there are features of our ability and disposition to make moral judgments a certain way that are quite similar to linguistic abilities and dispositions.<sup>11</sup> While this moral grammar may underlie moral judgment and action, moral analysis would attempt to discover the nature of the “grammar” itself. Just as there can be experts in judgment and action, there can be experts in moral analysis—moral philosophers are examples of this sort of expert, just as lin-

<sup>8</sup> This is, of course, controversial. See, for example, Jesse Prinz’s work on innateness; such as “Is Morality Innate?” in *Moral Psychology, Vol. 1*, edited by Walter Sinnott-Armstrong (Cambridge, MA: MIT Press, 2008): 367–406.

<sup>9</sup> John Mikhail, “Moral Grammar: Theory, Evidence, and the Future,” *Trends in Cognitive Sciences* 11 (2007): 143–52. Mikhail also discusses this in depth in *Elements of Moral Cognition* (New York: Cambridge University Press, 2011).

<sup>10</sup> *Ibid.*

<sup>11</sup> John Rawls famously developed the analogy between morality and language in his *Theory of Justice* (Cambridge, MA: Harvard University Press, 1971). Other writers, besides Mikhail, who develop the comparison are Erica Roedder and Gilbert Harman, “Moral Grammar,” (unpublished manuscript); and Susan Dwyer, Bryce Huebner, and Marc Hauser, “The Linguistic Analogy: Motivations, Results, and Speculations,” *Topics in Cognitive Science* 2 (2010): 486–510. It is worth noting that even if there is no UMG, the analogy can still be productive. The point of raising UMG is simply to provide a plausible account of why we might have underlying dispositions to think in moral terms.

guists are experts in linguistic analysis. Just as with language, there are expert users of the system and expert analysts, and those groups do not coincide, though they do overlap.

One appeal of the comparison between morality and language is that it provides a psychologically realistic understanding of how competent users of moral concepts deploy them without being able to articulate the norms that govern their deployment. Just as a person can be a fully competent speaker of English without being able to articulate English rules of grammar, a person can be a fully competent moral judge without being able to articulate the moral rules or norms that guide the judgment to which the agent is actually responsive in action. The judgment expertise is distinguished from the analysis expertise. An expert in linguistic analysis is better able to identify the norms. This helps to account for the fact that some disagreements over moral expertise are not genuine disagreements because writers are discussing different sorts of expertise. As Singer notes, there is a skill that moral philosophers possess that ordinary moral evaluators do not: moral philosophers are better able to identify the factors that people respond to as relevant—indeed, the factors that actually are relevant—in making moral evaluations themselves.<sup>12</sup> But it is a separate issue whether this skill translates into better evaluation. That will turn out to be an empirical issue, and skeptics about moral expertise often think that it does not yield better evaluation; but, if true, that view is not a skepticism about moral expertise per se, but rather a skepticism that expertise in analysis translates into expertise in judgment.

Often, in practice, the characteristics of expertise are interrelated. Consider automaticity and how it figures into expertise. To say that a judgment or an action is “automatic” is not to say that it is reflexive. This is important. Automatic behavior, in the sense used here, is norm governed, it is simply not deliberative. It is not like blinking, for example. Experts tend to make decisions very quickly, with little if any deliberation. The chicken sexer is a classic example of this. Some people are very good at determining the sex of chicks, though they are quite inarticulate when trying to describe what they do and often just can’t.<sup>13</sup> Thus, whatever criteria they are responding to are ones that they do not seem to have deliberated about at all. However, some qualification on this is required. Research in problem solving in science shows that experts do tend to deliberate a bit about what cognitive tools to bring to bear on a particular problem, but once they have decided *that* issue the problem solving is fast and reliable, relative to novices; novices, on the other hand, might quickly pick a tool, but then be stuck with an inappropriate one for the given problem, take much longer to actually solve the problem, and, further, do so less reliably. Consider problem

<sup>12</sup> Singer, “Moral Experts,” 117.

<sup>13</sup> David Eagleman, *Incognito* (New York: Random House, 2011), 57 ff.

solving in physics. This will require a person to be able to set up the problem, or represent it, properly.

Physics problems are often presented in terms of informal real-world objects and relationships among them. An important task of a problem solver is to obtain a formal representation of the problem by interpreting these objects and relationships into formal, abstract models of physics. Competence in solving a physics problem derives mainly from the skill of setting up the problem in terms of these models rather than from mere mathematical skill. An expert in physics is believed to possess a finite set of such models that may be applied to solve a large number of complex problems.<sup>14</sup>

Once that preparatory work is done, the expert proceeds quickly in solving the problem. So the expert is good at picking out the right way to model or understand a problem, the novice is poor at this, and this greatly slows the novice down when it comes to solving the problem relative to the expert. The expert is still characterized as, at some critical juncture in the problem-solving process, working much more quickly and reliably than the novice.<sup>15</sup> The process described above, in physics, describes the process of expertise in analysis, and holds for moral expertise of this sort as well.

One hypothesis as to why this is the case has to do with their greater experience, which rules out various options that the novice takes the time to consider. To use a philosophical term from virtue ethics, experts have learned what is to be “silenced” in deliberation in various contexts. This produces enormous gains in efficiency. Some psychologists actually take this automaticity model as empirical evidence against classical views of the moral agent as an explicit deliberator. Some even blame philosophy for misleading psychology about the nature of moral agency:

In our view, moral psychology is better served by jettisoning starting points that are motivated more by philosophical than by psychological considerations. . . . Rather than a “moralized psychology” whose parameters and terms of reference are set by certain philosophical goals . . . we opt instead for a “psychologized morality” that attempts to study moral functioning within the framework of contemporary psychological theories and methods.<sup>16</sup>

<sup>14</sup> Hyung Joon Kook and Gordon S. Novak, Jr., “Representation of Models for Expert Problem Solving in Physics,” *IEEE Transactions on Knowledge and Data Engineering* 3, no. 1 (1991): 48–54.

<sup>15</sup> *Ibid.* They also discuss ways in which machine problem solvers can also deploy models.

<sup>16</sup> Narvaez and Lapsley, “The Psychological Foundations of Everyday Morality and Moral Expertise,” 143.

The problem has been the focus on deliberation in moral agency, they argue. Empirical research shows that deliberation plays very little role in a great deal of behavior that we think, intuitively, is subject to moral evaluation. Thus, they argue, this old model, influenced by philosophical goals that try to tie moral behavior to deliberation, is mistaken. This is quite relevant to the issue of expertise because Narvaez and Lapsley are attributing a false model of agency—and corresponding agency expertise—to philosophical accounts.

Thus, it is worth pointing out that *mere* automaticity is not a worry for the classical view at all. If Narvaez and Lapsley's view reflects the considered opinion of empirical moral psychologists, then they are attacking a "strawman." Aristotle's account of moral virtue has generally been taken to be the most intellectualist, requiring practical reason for virtue, and yet on the most common interpretation of Aristotle's account of moral virtue, automatic behavior is compatible with virtue as long as the automatic behavior was the result of the right kind of "training up" of the agent. Indeed, this would be how one understands the method by which the regulation of the behavior occurs. The agent is, in the past, made aware of the relevant reasons and actually uses them to guide behavior. The appropriate training involves deliberation *at some point*. While I believe that even this watered down commitment to deliberation is incorrect, the evidence cited by Narvaez and Lapsley is not in itself harmful to the Aristotelian picture.<sup>17</sup> Indeed, some of what they write seems to indirectly support it, as when they note that automaticity can develop via habituation of routine tasks. It is almost a truism in normative ethics today that being a morally good person does not require that one engage in constant deliberation. But even much earlier theorists such as Jeremy Bentham never required that the good agent constantly engage the Utilitarian calculus—the efficiency costs are much too high.<sup>18</sup> Indeed, one mark of the expert is the reduction in efficiency costs in achieving a better-than-normal result.

Narvaez and Lapsley also argue that the "principle of phenomenalism," a holdover from moralistic moral psychology, has misled philosophers. This principle maintains that "... a behavior has no particular moral status *unless it is motivated by an explicit moral judgement.*"<sup>19</sup> This puts huge weight on "the subjective intention of the moral agent."<sup>20</sup> They

<sup>17</sup> See my "Dubious Virtue Psychology" in *Uneasy Virtue* (New York: Cambridge University Press, 2001), 42–62.

<sup>18</sup> Bentham makes this quite explicit in *The Principles of Morals and Legislation* (New York: Hafner Publishing Co., 1948), 31, when, after describing the many different parameters along which pleasure and pain are measured to determine the moral correctness of an action or policy he states: "It is not to be expected that this process should be strictly pursued previously to every moral judgment, or to every legislative or judicial operation."

<sup>19</sup> Narvaez and Lapsley, "The Psychological Foundations of Everyday Morality and Moral Expertise," 141.

<sup>20</sup> *Ibid.*

argue that this model is defective on a number of fronts, but one major problem is that it is incompatible with expertise, since expertise involves automaticity which is, on their view, clearly nondeliberative (indeed, it is understood in contrast to the deliberative). One issue the discussion raises is the implicit identification of the moral expert with the morally good person. I believe a distinction needs to be made between the two, though it seems plausible to hold that the morally good person is one kind of moral expert, an expert practitioner, of sorts. But there is also the sort of expert in judgment, and (at least on the view that I favor), judgment and practice can come apart. But with this distinction in mind, each type of expertise can be marked by the same conditions. For example, the expert practitioner acts “automatically” and the expert judge judges “automatically.” There may be a very strong correlation between the two, but they can come apart.

I agree with Narvaez and Lapsley that expertise is often marked by automaticity, but disagree that this is necessarily a problem for the deliberation model (though I do think that there are plenty of *other* problems for this model). However, the automaticity model allows us to understand moral expertise in ways that have been philosophically controversial. So, not only are there moral experts, but there are lots of moral experts if we understand an expert to be good at what she does *relative to a specific contrast*. What is rarer is metamoral expertise—that is, expertise about morality itself (which itself is distinguished from expertise about moral expertise). And those who argue against moral expertise have in mind a person who is expert relative to the contrast of the “average Joe.” That view is that there are no experts once a certain level of competence is reached. However, this again seems implausible. It is certainly implausible when it comes to experts in analysis. What about expert judges and practitioners? It is important to point out that this kind of skepticism is different from Ryle’s. There is a kind of skepticism about expertise that grants moral knowledge, but holds that once moral competence is achieved, there are no people that know more than anyone else. This line of thought can be traced, I believe, back to views about autonomy as the source of normativity.<sup>21</sup> We are all equal as autonomous beings, and possess the same autonomous capacities. But this line of thought confuses two issues: one is the issue of equality of capacity, and the other is the issue of the development of that capacity. Even if we agree that all competent adults possess the same capacity for making moral determinations, it does not follow that the capacity is exercised equally well. We can dip into the language analogy again and note the competence/performance distinction as well. Even if (and this seems like a very big “if” to me) all rational competent adults were equal in terms of their moral competence, there could still be significant variation in performance. Performance defects

<sup>21</sup> I discuss this in “Autonomy and the Asymmetry Problem for Moral Expertise.”

can be intentional or unintentional, and may not reflect on basic competence. A person who, under stress, misspeaks, is not linguistically incompetent. Similarly, a person who, under stress, slips up and accidentally does something bad, is not morally incompetent.

One feature that distinguishes moral expertise is that the moral expert can distinguish better from worse moral judgments (relative to the naïve sorter). Whereas a young child might think that it is morally permissible to hit Johnny for taking his cookie, the more mature judge would disagree. This raises an epistemic issue that has concerned some writers on moral expertise: How do we know what judgments are the better ones?<sup>22</sup> If we already know, then what distinguishes the expert from the non-expert? This does not seem to me a serious problem for the notion. This can be handled the same way we handle any kind of normative expertise. Suppose that I go to the art museum with my friend Sasha who is an aesthetic expert. My naïve impression upon viewing an Ancient Egyptian tomb painting is that it is not very good. Sasha, however, disagrees and is able to explain why the painting is in fact quite remarkable. She understands the rules and conventions that governed the Ancient Egyptian artists, and their rejection of what to us would appear as proper perspective. To the extent that I find the explanation convincing, then I will view her ability to sort aesthetic judgments as superior to mine. Note that this does not commit one to the view that the expert *must* be able to explain the judgment. That is an additional skill. I want to allow for the possibility of inarticulate sorters.

Again, the analogy I would like to pursue here is with linguistic expertise. Moral facility develops in children as they interact with others, the same with language. But linguistic expertise can itself come in a variety of forms. There is facility in speaking a language and facility in reading a language, and these facilities are quite distinct. It is much easier to learn to speak than to read. If we pursue this analogy, one question is whether morality (i.e., making moral judgments) is like speaking or reading.

Speaking well follows patterns, though the speaker may not be able to articulate the patterns. However, when one is learning a language later in life, one does need to engage in a good deal of deliberative practice in order to generate the same level of automaticity one has achieved in one's native language. Further, it is quite possible to forget a language. What makes Ryle's observation so plausible is that moral norms are so overwhelmingly significant in our interactions with others that we are not going to fall out of practice or familiarity with them.

One potential problem with this approach is the existence of moral disagreement. We might think of experts as those who get morality right, relative to the naïve practitioner. But if we are all experts of a sort, then

<sup>22</sup> See Brad Hooker's "Moral Expertise" in *Routledge Encyclopedia of Philosophy*, ed. Edward Craig (New York: Routledge, 1998): 510.

why is there such widespread disagreement? One could hold, as those who argue against empirical moral relativism hold, that there really is not much in the way of actual moral disagreement, that is, disagreement about values. Rather, apparent moral disagreement is due to disagreement about matters such as long-term consequences of available options. Further, experts do disagree. The existence of “expert” disagreement itself does not tell against moral expertise any more than it does against any other sort of expertise.

But there are other options as well. It could even be that experts are mistaken *systematically* about some issue—it does not seem to be a feature of our concept of expertise that experts are infallible. They just *tend* to get it right relative to naïve practitioners. Further, some linguistic mistakes can be disguised inasmuch as those mistakes are coordinated. This will hold for moral mistakes as well. Consider “depth charge” sentences, such as:

- (1) No head injury is too trivial to be ignored.<sup>23</sup>

People generally take (1) to mean that “all head injuries, no matter how small, should be examined and treated.” But, as linguists note, this is ruled out by the grammar of the sentence. This sentence is structurally just like:

- (2) No missile is too small to be banned.

This sentence means that “all missiles no matter how small should be banned.” Thus, (1) actually means the exact opposite of what most speakers take it to mean! The speakers then, the linguistic experts, get it *systematically* wrong. Not only systematically wrong, but in a coordinated way—as when the members of a group of people rely on, and coordinate their behavior according to, a mistaken timepiece.

By analogy, this can be the case for moral expertise as well. Moral experts arrive at correct moral judgments reliably, relative to naïve practitioners. But this is compatible with systematic errors on their part. Perhaps the belief that “killing is intrinsically worse than letting die” is an example of systematic error. How is such error to be diagnosed? The same way as linguistic error—*meta* field experts. Meta-linguistic experts are good at figuring out linguistic errors, metaethics experts may be as well, for moral judgments.<sup>24</sup> Ordinary experts do not typically have

<sup>23</sup> Anthony J. Sanford and Patrick Sturt, “Depth Processing in Language Comprehension: Not Noticing the Evidence,” *Trends in Cognitive Science* 6 (2002): 384.

<sup>24</sup> More distinctions are possible. There is the expert actor, the expert deliberator, the expert advisor, as well as the metamoral expert.

metaexpertise in their field (for example, good baseball players might not make the best coaches). In the linguistic case, the expert can detect fairly reliably what is grammatical and what is not. But again, this will not be infallible and there will be systematic errors that the meta-linguistic expert is in a position to identify. Consider the case of “garden path” sentences such as:

(3) The prime number few.<sup>25</sup>

Most people view (3) as ungrammatical, but they are wrong. Consider

(4) The mediocre are many but the prime number few.

People with great, automatic, facility with language will tend to make these mistakes. It is up to the meta-linguistic expert to recognize and diagnose the mistake.

We need to distinguish “grammatical” from “acceptable.” (3) is grammatical (*vis à vis* actual word usage), though not acceptable in that it fails to convey reliably what the speaker wants it to convey. An example of an expression that is not grammatical, but acceptable:

(5) a not unhappy person.<sup>26</sup>

A grammatical expression is normatively correct relative to the norms of grammar. An acceptable expression is one that normal, competent, speakers of the language would assent to, in the sense of holding it to be a proper sentence in the language (i.e., they assent to its grammaticality). The analogy with morality will preserve such a distinction between what is morally permissible given the norms of morality and what a normal, competent, user of moral expressions would assent to as deriving from those norms of morality.<sup>27</sup>

Consider as an analogy

(6) Killing is worse than letting die.

<sup>25</sup> Steven Pinker, *The Language Instinct* (New York: Harper Perennial, 1998).

<sup>26</sup> D. Terence Langendoen, “On a Class of Not Ungrammatical Constructions,” *Journal of Linguistics* 18, no. 1 (March 1982), 107. Other possibilities are: “. . . just between you and I. . .”; “There are frogs in the pond.” These are discussed in Nicolas Sobin, “An Acceptable Ungrammatical Construction,” in *The Reality of Linguistic Rules*, ed. Susan D. Lima, et al. (Philadelphia: John Benjamin, 1994): 51–65.

<sup>27</sup> Note that this is itself distinct from the issue of truth. If Mackie is right then sentences expressing moral claims can be well formed even though all of them are false. What is important is simply that there is some standard of assessment. If not “true/false,” then “appropriate/inappropriate” would be fine.

Suppose that moral norms are utilitarian. (6) is false relative to those norms, and thus (6), strictly, expresses a claim that does not follow from the norms of utilitarianism. Is it acceptable? That would depend on its usage, but a case can be made for the acceptability of (6). It seems true that normal, competent, users of moral terms would assent to (6) as following from moral norms since they are being influenced (probably) by some form of representative heuristic. Since the typical killing seems ill intentioned and the typical letting-die does not, the normal user might well assent to (6) even though (6) does not follow from utilitarian norms.<sup>28</sup> The moral expert may get at moral truth as well, though this depends on the truth of certain meta-ethical views. On the view that I recommend here, there need not be truth in morality for there to be moral experts.

What of the other conditions psychologists have associated with expertise in general? The first I noted was the ability to make fine and nuanced discriminations. This has long been noted as crucial to “critical” judgment in philosophy. In David Hume’s “Of the Standard of Taste,” much is made of the aesthetic expert’s ability to distinguish the slight taste of iron and leather in a cask of wine.<sup>29</sup> Analogues in normative ethics include the ability to tease apart what is operating on our moral judgments, especially in cases of normative ambivalence: for example, if Melinda must assassinate an evil dictator using just a ballpoint pen, reactions might be mixed. The nuanced judgment, however, is that she did the right thing, though we might be concerned about her dispositions and overall character. Here we leave aside the issue of Melinda and whether she is a moral expert—it is the evaluator making the nuanced judgment who is the expert.

The related claim, that because of their ability to make nuanced discriminations, experts actually see the world differently, is a bit more interesting and controversial.

... [E]xperts have large, rich, organized networks of representations (schemas) containing a great deal of declarative knowledge about the domain, and well-organized, higher interconnected units of knowledge in the domain. . . . Because they have more and better organized knowledge in a domain, experts perceive things differently than do novices. . . . Experts in morality, like experts of all kinds, can be expected to perceive and act upon the world in a markedly different way than do moral novices. For example, experts in moral sensitivity are able to more easily pick up on the morally relevant affordances in the environment (e.g. What is my role in

<sup>28</sup> Note that the particular form of the example is not important. If one does not like this one, or agree with its use, just pick another.

<sup>29</sup> David Hume, *Selected Essays*, edited by Stephen Copley and Andrew Edgar (New York: Oxford University Press, 1998): 133–53.

this situation? What should I do? What am I capable of doing? What does the context allow?)<sup>30</sup>

Again, we need to note a distinction between the expert who knows what to do, and then the expert who knows what to do and does it. These two may come apart, on my view. Otherwise, this passage reads remarkably like some of John McDowell's work on virtue, except that the virtuous person is read as a moral expert.<sup>31</sup> McDowell holds that the morally virtuous person perceives what is morally significant and is properly motivated by that perception.

What does seem to be the case, however, is that even the nonvirtuous expert sees the world differently from the novice in a way that involves having different cognitive states about the world, and this would be another point of disagreement with Ryle.

Again, Ryle argued that there can be no moral experts because there is nothing for them to be expert *about*, there is no moral knowledge. This is because we cannot forget morality, in the sense that we do not forget what is right or wrong once we have learned it. However, this datum does not seem entirely accurate. Suppose that Alex has decided to become a vegan, for moral reasons. He has learned that eating dairy products is wrong due to the pain and suffering caused by the factory farming procedures used in most dairies. But veganism is not something that is, where he lives, commonplace. He finds himself forgetting that it is wrong to eat cheese when he goes out with his friends. Admittedly, this kind of forgetting will not be common, but it will not be common primarily because there is widespread agreement about the basics of morality. What makes the above case plausible is that veganism is not *generally* considered a moral issue in the present culture. Social reinforcement is lacking, whereas it is not lacking for bits of learning that we grow up with, such as "causing pain is bad," and so forth.

If there is some additional cognitive element to moral expertise that helps explain how experts perceive things differently, then how can we understand this? One route is to argue that we can use the distinction between seeing and *epistemic* seeing, though understood more broadly in discussing moral perception. This is a distinction made famous by Fred Dretske in *Seeing and Knowing*.<sup>32</sup> Just plain seeing is something we clearly share with other animals—it is "devoid of positive belief content." Epistemic seeing, on the other hand, is not. Further, one can view epistemic seeing in two ways—as successful and as unsuccessful, depending on how one views the "that" in "seeing that x." When Jane walks into a

<sup>30</sup> Narvaez and Lapsley, "The Psychological Foundations of Everyday Morality and Moral Expertise," 150–51.

<sup>31</sup> John McDowell, "Virtue and Reason," *The Monist* 62 (1979): 331–50.

<sup>32</sup> Fred Dretske, *Seeing and Knowing* (Chicago: University of Chicago Press, 1988).

crowded room, she sees people, but she does not see that there are seventeen people in the room. Suppose she believes that there are fifteen people in the room, and would assent to the false claim that “Jane sees that there are fifteen people in the room.”? Is that “seeing that,” or not? For our purposes we can just note the two ways of understanding the “that” and also note that in discussing the issue of expertise, the expert tends to “see that” successfully in his or her domain of expertise.

Suppose that Sophia is a zoologist and she takes her daughter Andrea with her on a specimen collecting expedition in the Mojave Desert. Andrea, walking through the desert, sees lots of sand, she sees patterns and swirls in the sand, and she sees various plants she knows to be cacti. But she misses things that her mother—who sees exactly what Andrea sees—does not miss. Her mother sees that a particular swirl in the sand is a rattlesnake mark. Sophia also sees that a bird nesting in a particular kind of cactus is a rare species. Thus, her mother is able to make finer discriminations due to her expert knowledge, and this in turn really does influence how she sees the world around her. Given a certain set of independent motivations, too, it partly explains her actions. Since she has a desire to collect rare specimens, it explains her action of taking the plant with her. If the collection is warranted, it also justifies taking that plant (rather than some other plant).

In the moral case the moral expert, let’s say, picks up on features of the situation that are morally relevant. Quite often these features will be obvious to the normal competent adult. In hard cases, this is not true, and this is where we get an upper tier of expertise. Normal competent adults may not be familiar with situations in which they need to decide whether, for example, to keep someone alive using extraordinary means or not.

One final feature of expertise discussed by psychologists concerns the skill sets they develop in accessing and applying knowledge of past cases. This suggests that experts are skilled at analogical *retrieval* and *mapping*. There is a good deal of evidence in psychology that some people, for whatever reason, are much better at these tasks than others. However, one factor that influences a person’s ability to retrieve the appropriate analogue to compare with the target problem is expertise in a particular area. Experts retrieve relevant analogues much more reliably and much faster than novices, and this increases the speed and accuracy of their analogical reasoning.

There will be people who are better able to develop analogies that are illustrative, and allow for extensions of moral knowledge. These will be experts relative to the class or normal adult practitioners. Such people would be able to extend accurate moral judgment to novel cases—the sorts of cases that often turn up in biomedical ethics courses as “hard” cases. Thus, the account suggested in this paper is an account of moral expertise that, like linguistic expertise, is layered according to contrasts with various classes of naïve practitioners.

However, one problem is that the analogy does not seem to reflect a crucial distinction for moral judgment—that is the distinction between convention and morality. Language is conventional and morality is not conventional. Moral norms possess the sort of necessity that linguistic norms do not. Linguistic norms are arbitrary in a way that moral norms are not. That one ought never split the infinitive is a convention of language. That one reads from left to right is a convention of the English language. And morality is not like this. The fact that something is *fundamentally* “right” or “wrong” is not a matter of convention. Of course, the criticism goes, there is a superficial way in which things can be wrong via convention. In a society in which the convention is for people to drive on the right side of the road, it would be wrong to drive on the left side. But what makes it wrong is not itself conventional. What makes it wrong is that in defying the convention one is risking significant harm to others. And that this is wrong is not a matter of convention.

There are two ways to respond to this line of argument. The first is to note that it does not undercut the analogy, it simply limits it. No analogy will be perfect, and that is actually a crucial feature of the analogy itself. The analogy is limited to pointing out that, like linguistic expertise, moral expertise of a certain sort can exist without expertise in analysis, which is a different sort of skill. This is because a person can be a good moral practitioner, actor, or judge, even absent an articulate understanding of the norms that govern the practice. Language gives us another way of seeing how this works.

The second is more aggressive. The second strategy is to note that whether or not one sees this as a disanalogy depends on how one views morality. Is it like logic? Or, is it more like aesthetics? Both logic and aesthetics are norm governed, though the norms of logic are not considered conventional at all, whereas this seems more likely in the case of aesthetics. I take it that the underlying worry about the disanalogy is that there is something about morality such that moral truth is not contingent on anything other than what the moral facts—stance-independent moral facts (in the sense of not depending upon facts about human nature, for example)—happen to be. But there are no “stance-independent” linguistic facts in the sense of being independent of features of human nature. Answering this sort of question is beyond the scope of this particular essay. However, if one decides to pursue the analogy beyond the limited aims stated earlier, then the most promising way to do it would be to show that truths of morality are contingent on facts of human nature—for example, by adopting a constructivist approach to morality. This would undercut the criticism that morality and language are so dissimilar that the analogy is not a useful one.