

## ***A Plea for Moral Deference***

Gopal Sreenivasan

Crown Professor of Ethics

Duke University

Philosophy, Arts & Sciences

[gopal.sreenivasan@duke.edu](mailto:gopal.sreenivasan@duke.edu)

### **ABSTRACT**

It seems to be a commonplace of the philosophical literature that there is no such thing as moral expertise. Or perhaps, more narrowly, that there is no such thing as justified deference to moral expertise, when there is moral expertise. On the other hand, a warrant for moral deference seems to have a secure place in everyday moral experience. It is illustrated, for example, by the ubiquitous phenomenon of taking moral advice (this includes a role for exemplars of individual moral virtues, but is not limited to exemplars of virtue). In this paper, I shall defend moral deference against overblown philosophical skepticism. I hope to contribute to rehabilitating the notion for some role in moral theory.

### **KEYWORDS**

Moral deference, moral expertise, moral virtue, ethics, virtue ethics.

Imagine that you face a practical situation in which you do not know what the morally right thing to do is. Suppose, moreover, that you also happen to know someone who does know what the right thing for you to do is (and she is in a position to advise you). Should you accept her moral advice? In due course we shall encounter and discuss some examples that fit this bill. It will also be useful to refine the parameters of our question in a little more detail along the way. But we can make a reasonable start in fairly general terms.

In the recent philosophical literature around this question, two facts stand out as fixed points of the discussion, with one of them having greater salience than the other. The first fact is that the question has some importance for moral theory, especially perhaps for theories of virtue in particular. The more salient fact is that, sociologically, there is a wide consensus among participating philosophers that, in some important sense, the correct answer to the question is ‘no, you should not accept the other person’s moral advice.’ Naturally, my two qualifications call for some elaboration.

A narrower way of construing the relevant consensus would be to express the point of agreement as the proposition that accepting the moral advice is somehow

notably objectionable: it carries some kind of toxic stain.<sup>1</sup> While some parties to this agreement hold that it is nevertheless sometimes permissible to accept the advice—that, under some conditions, the objection may be overcome—all parties agree that there is something to be overcome. Indeed, a common project in this literature, undertaken even by ‘defenders’ of moral deference or testimony,<sup>2</sup> is to *explain* the stain that attaches to accepting moral advice, i.e. to identify or diagnose it properly. The presence of the stain accounts for the ‘important sense’ in which you should *not* accept the advice, even if, all things considered, accepting it remains permissible.

My own view is that no stain attaches to your accepting the advice of someone who knows (when you do not) what the right thing for you to do is. You may accept such advice, and often you should accept it. More significantly, in the basic and most instructive version of the case, there is no good objection to your accepting it. In arguing for this conclusion, I not only decline to join the recent philosophical consensus, but I reject it. (Of course, I do not deny that the consensus itself exists—hence, ‘sociologically’). Unlike many of its other defenders, then, the plea I shall be making for moral deference will be a plea without excuses.

Now different contributors to this debate enter it along different terminological pathways. The relevant terrain is defined by a nexus of inter-relations among the terms, ‘moral expertise,’ ‘moral testimony,’ and ‘moral deference,’ where (roughly) non-experts *defer* to *experts* in relation to their *testimony*. It is possible to distinguish sharply between any pair of these terms, thereby severing one term from the nexus. The motivation for so doing is usually to establish its innocence by disassociation. But, whatever their motivation, contributors often leave at least one of these terms outside the scope of their enquiry altogether. I shall do the same.

In what follows, I pay no particular attention to moral ‘testimony.’ Insofar as there is any reason for this, it is because I shall also be ignoring a subsidiary debate that commonly arises here concerning an alleged asymmetry between the moral and non-moral cases.<sup>3</sup> Everyone (or almost everyone) admits that there is nothing objectionable about deferring to non-moral testimony. Given the background consensus that something is wrong with deferring to moral testimony, a further question therefore arises of how to explain the resultant asymmetry. Among other things, framing this asymmetry in terms of ‘testimony’ facilitates comparisons with an established epistemological literature on (non-moral) testimony. Clearly, I reject

---

<sup>1</sup> For example, it is ‘unacceptable’ (Hopkins 2007; Hills 2009), ‘off-putting’ (McGrath 2011), or ‘fishy’ (Enoch 2014).

<sup>2</sup> See, e.g., Sliwa (2012) and Enoch (2014).

<sup>3</sup> See, e.g., Driver (2006), McGrath (2011), and Howell (2014).

the presupposition of the comparison. To some extent, however, I am also simply opting for a narrower scope.

While I ultimately wish to focus on evaluations of moral *deference*, I shall begin with moral expertise. For the most basic objection to moral deference is that there is simply no such thing as moral expertise, and hence nobody to whom one might defer morally. Even though I believe that this objection is both mistaken and confused, it remains well worth discussing. I shall then introduce an example and develop my defence of moral deference in relation to it. Throughout I shall argue in terms of morality quite generally. Yet I am also interested in the ramifications for deference to exemplars of the moral virtues, specifically. Accordingly, I shall close with a coda on virtue. In the coda, I shall argue that, in several respects, the case for deference within the province of virtue is easier to make than the generic case for moral deference.

### 1. *Moral expertise*

Let us begin by considering an ‘in principle’ version of the objection that there are no moral experts. So construed, the objection lends itself to either a negative or a positive formulation. In its negative formulation, the objection derives from a denial that there can be any in principle moral epistemic elite—a denial, in other words, that there is any sub-class of individuals who know (or are even able to know) moral truths that the rest of us cannot know. In its positive formulation, it derives from an affirmation that, in principle, moral truths are all fully accessible to everyone, i.e. to every ordinary person. Although the historical origins of this view are commonly—in the philosophical literature, anyhow—attributed to Kant, they really belong to a cultural legacy to which Kant himself is much more heir than testator, namely, the legacy of the Protestant Reformation. Fundamentally, this version of the objection originated in an anti-clerical critique: specifically, in the idea that there are no human gate-keepers to salvation.<sup>4</sup>

In any case, as a matter of substance rather than history, the basic idea inspiring the objection is plainly very appealing, since it amounts to the democratization of morality itself. Equivalently, it applies something like a principle of equality of opportunity to the achievement of moral knowledge, and thereby to the achievement of a morally good life. Unfortunately, however, just because an idea is very nice—noble, even—that does not make it true. On inspection, moreover, this particular idea turns out to be philosophically suspect. Let us examine the matter under its positive formulation. What might account for the fact, if it is one, that the truths of morality are all fully accessible to every actual human being?

---

<sup>4</sup> See, e.g., the helpful account in Taylor (1989, ch. 13).

At least when human epistemic capacities are taken as a given, the explanation for the universal accessibility of moral truths arguably has to rest *either* on a giant coincidence *or* on the operation of some prior constraint on the moral truths themselves. That is to say, in the latter case, that this accessibility results from morality's having been 'bent to fit' our epistemic capacities. Which particular epistemic capacities served as the target for this bending depends on whether every actual human being is equally capable (in principle) of working out the requirements of morality. In the more difficult—but presumably, much more plausible—case in which actual human beings (even 'normal' ones) are *not* equally capable in this respect, the prior constraint would have the effect of bending morality to fit the lowest common denominator among (normal) human epistemic capacities.

For simplicity, let me reject the giant coincidence out of hand.<sup>5</sup> What remains is the possibility of explaining the universal accessibility of moral truths on the 'bent to fit' model. But this explanation seems to commit us to some kind of constructivism about morality. As far as secular accounts of morality are concerned, then, the explanatory basis of the present objection to moral expertise appears to be inconsistent with the objectivity of morality.<sup>6</sup> Whether or not that counts as a

---

<sup>5</sup> How demanding a coincidence is required depends, inter alia, on whether the content of morality is fixed in advance of the comparison or not. On one model—probably more appropriate to the historical view, but much less philosophically robust—we assume a particular content for morality (conventional Protestantism, say) and then compare the epistemic capacities required to grasp *that* content to our actual (lowest common denominator) capacities. The claim on offer here is that the terms of this comparison happen to coincide. On another model, moral truths are held to be universally accessible, *whatever* the content morality turns out (correctly) to have. While this is philosophically more robust, the coincidence it requires is also correspondingly more demanding: for here the claim is that, of all the manifold contents morality might have, those contents that exceed the reach of our actual (lowest common denominator) epistemic capacities just happen, *all* of them, to be false. Of course, the staggering extent of this coincidence can always be avoided by falling back on the first, less demanding model. But the consequent reduction in coincidence will be encumbered by the presupposition that the particular content nominated for morality is both correct and complete. Either way, the position seems unsatisfactory.

<sup>6</sup> On theological premisses, by contrast, it is still possible to reconcile a version of constructivism about morality with its objectivity. Moreover, if the bending is to be God's work, it seems that the desired congruence between our actual epistemic capacities and the capacities required to grasp morality could in principle be achieved by bending in *either* direction. In other words, a further option of *bending our capacities* (to bring morality within their reach) then comes into view. While a secular account can, of course, also allow human epistemic capacities to change, I am assuming that it has no way to make sense of the idea of bending them to fit morality.

decisive strike against the objection, it is certainly philosophically problematic.<sup>7</sup> In any case, that is my ground for classifying the objection as a mistake.

I need not insist on this point, however, because the objection is also confused; and its confusion alone is adequate for my purposes. To see this, let us stipulate that every normal adult human being is equally capable (in principle) of working out the requirements of morality. Moral truths are therefore all fully accessible in principle to every ordinary person. We can regiment this landscape by saying that there are no moral ‘experts’ (in it). All the same, it simply does not follow that there is nobody to whom one might defer about what morality requires one to do.

To generate that further conclusion, even for a given point in time, we would need to be shown that every ordinary person had *developed* his or her in principle equal epistemic capacities *equally*, i.e. to the same extent as everyone else (at that time). Alternatively, differential moral learning among some group of individuals is both consistent with their underlying epistemic capacities all being equal in principle and yet inconsistent with there being nobody in the group to whom any one of them might defer on moral questions.

Those who have learned more are natural candidates for the role of someone to whom others who have learned less might defer. It remains a separate question, of course, whether there is anything objectionable about moral deference (and we shall join that question below). However, the present issue is not whether moral deference is objectionable, but rather whether it is even possible, in the minimal sense of there being anyone available to whom others might defer on moral questions. On the face of it, one person’s having some moral knowledge that another person lacks – or, perhaps better,<sup>8</sup> her *reliably* having such knowledge – is sufficient to put the first person in the role of someone to whom the second person might defer morally. To occupy this role, the first person need not be a moral ‘expert’ in any stronger sense than that,<sup>9</sup> and certainly need not be an expert in the sense of knowing some moral truths to which the second person lacks epistemic access in principle.

---

<sup>7</sup> Here as elsewhere, we may be reminded of the familiar symmetry between *modus ponens* and *modus tollens*. McGrath (2011), e.g., drives something like this argument in the other direction, taking the (asymmetrically) ‘off-putting’ character of moral deference as a ‘datum,’ and deriving a challenge for moral realism from it.

<sup>8</sup> We can refine the relevant conditions more fully in the next section.

<sup>9</sup> I have no stake in the label: it makes no substantive difference whether ‘expert’ is in any sense an appropriate label for the ‘object of deference’ role. If it is, then the confusion in the objection we are discussing can be diagnosed as equivocation between weaker and stronger senses of ‘moral expert.’ If ‘expert’ is a wholly inappropriate label, then the objection’s confusion is to have (wrongly) assumed that the existence of moral experts is a condition of the possibility of moral deference.

Logically, then, the universal accessibility of moral truths (even when granted for free) is hopeless as a basis for pre-empting the possibility of moral deference. Nevertheless, it should also be acknowledged that ‘differential moral learning’ is not merely a logically coherent spanner that happens to lend itself to insertion in the anti-deference crusader’s works. On the contrary, I take it that differential moral learning is a plain fact of ordinary moral experience, and indeed a massively common one. All of us, presumably, are acquainted with people who are better than we are, or more reliably knowledgeable anyhow, in some or other department of morality (if not many such departments). There is no shame in admitting this. Nor (I’m guessing) does it require a lot of reflection, but only a little honesty or humility, to recognize oneself in that portrait.

One could always try to factor differential moral learning into its causal constituents. Besides the omnipresent possibility of differential underlying epistemic capacity (itself a complex of factors, clearly), there is certainly differential *experience* as well, and no doubt more. I shall not pursue this sort of analysis, since the raw fact of differential moral learning is adequate for our purposes. Still, in case some resist acknowledging the raw fact in the first place, it may be worth briefly going one more round.

Differential experience can be factored into differential net quality and differential quantity. Opening an explicit place for quantity reminds us that the ‘wise people’ who populate legend and folklore are invariably wise *old* people. That is to say, brute relative age—certainly, a generation gap—makes a prime contribution to differential moral learning,<sup>10</sup> which is relevant here because real people are always distributed across a generational spectrum, unlike the weightless contemporaries of abstract analysis. On reflection, this is itself enough to yield our conclusion.

Consider, e.g., the platitude that it is appropriate (compulsory, really) for children to defer morally to their parents or elders. What makes a child’s moral deference to his elders possible is their differential moral learning. However, as long as this differential persists into the child’s adulthood, the upshot will precisely be differential learning between adults (the grown child and his elders still), i.e. our raw fact. To resist this conclusion, one would have to deny that this differential typically does persist for a significant interval, which requires one to embrace one of the following nettles: Either there is some early point in adult life at which moral learning effectively ceases or else the bare attainment of majority obliterates any

---

<sup>10</sup> Its contribution is consistent, of course, with some cases of precocious moral wisdom, as well as with the age difference between two people’s declining in epistemic significance over time. I insert the qualification ‘typically’ into my claim about the persistence of differential moral learning, in the next paragraph, to cover just these points.

remaining gaps in the former child's moral knowledge (relative to his elders). While the latter proposition is scarcely credible, the former is not at all plausible either.

## 2. *Moral deference*

So let us take it that moral deference is possible and return to the question of whether it is objectionable. To focus our discussion, consider the following everyday sort of example:

*Country cousin.* Suppose that my poor cousin from the country is coming to town for a few days to interview for a job. He has asked me to put him up during his visit, since a hotel bill would be a real hardship for him. But my apartment is very small, with hardly enough room for my immediate family, which includes a baby. Since there is no spare couch, I would have to let him sleep on the floor. That would work, albeit with some discomfort for all concerned, especially him. The alternatives seem to be footing his bill at a modest hotel nearby or finding some friend with a larger place to help me help my cousin. Each of the available options—floor, footing, or friend—has its disadvantages.<sup>11</sup> I have no hesitation, let us say, in agreeing to my cousin's request. But I also have no idea which of these ways of putting him up is best.

There are various ways in which I might consider how to respond to my cousin's request. I might simply wonder which of the options is best. Or I might wonder, of some particular option, whether *it* is required (or somehow, best). Or I might wonder something yet more specific. With this case, it seems natural to wonder whether *generosity*, e.g., requires me to take the footing option. To wonder, that is, whether offering my cousin the floor or my friend's couch would constitute a criticisable lack of generosity on my part.<sup>12</sup> Let me begin by following this particular thread.<sup>13</sup>

By hypothesis, I myself have no idea whether or not generosity requires me to take the footing option in *Country cousin*. (While I may have a nagging suspicion to that effect, I have no idea how to evaluate that suspicion). How then might I proceed, given my ignorance? One possibility, evidently, is to *ask* somebody who

---

<sup>11</sup> The hotel bill would be a strain for me, though I can certainly manage it better than my cousin. On the other hand, he might feel rather awkward or offended being sloughed off on someone he does not know.

<sup>12</sup> We need not concern ourselves with the issue of what the force of this criticism would be, exactly. We may simply suppose that I innocently aspire not to merit the criticism.

<sup>13</sup> In the coda on virtue, I shall return to comment on the relation between this narrow thread and the widest question, what is the right thing for me to do (in *Country cousin*)? I shall claim, though, that it makes no real difference to the status of moral deference.

does know what generosity requires (or, more generally, someone in a better epistemic position). However, in order for this quite general possibility to be practically available to me, two further conditions have to hold.

To begin with, trivially, the situation must leave me enough latitude to seek counsel (sometimes immediate action is required). Furthermore, someone who knows what generosity requires has not merely to exist, but to be *known* (and available) *to me*. There is some question about how this second condition gets to be satisfied, but for the moment I shall simply stipulate that it holds.<sup>14</sup> Let us say that I have a friend available, Gina, who is known to me as a model of generosity (i.e., I know that she is *reliably* knowledgeable about what generosity requires). For convenience, I shall refer to this version of the case, which builds in the two further conditions, as *Country cousin (plus)*.

Now suppose that I ask my friend Gina for advice and that she says there is no particular reason to foot my cousin's bill.<sup>15</sup> "You could," she says, "but it would be going over the top. Offering him the friend's couch is perfectly good. You can offer him your own floor as well, to make it clear that you are not trying to slough him off. He can choose." From here, we could proceed straight into an examination of whether there is anything objectionable about my accepting Gina's moral advice. But this may trample over a distinction on which some critics of moral deference have wished to insist, namely, a distinction between accepting moral advice and moral *deference* proper.<sup>16</sup>

As applied to *Country cousin (plus)*, this distinction turns, roughly, on whether I accept Gina's judgement simply on the basis of her known reliability as a model of generosity (deference) or whether Gina's role extends instead to bringing me somehow to grasp myself *why* generosity does not require me to foot my cousin's bill (advice). Thus, to clarify that our case offers an example precisely of 'deference,' so understood, let us specify that I accept Gina's judgement here simply on the basis of

---

<sup>14</sup> This stipulation is dialectically legitimate because (following on the argument in the previous section) the controversy has now shifted, in effect, to the question of whether moral deference is objectionable, *given that it is possible*. My stipulation merely serves to tighten the scope of the italicized concession. Nevertheless, a profitable question does remain to be discussed here, and I shall return to it (too) in the coda on virtue.

<sup>15</sup> To forestall irrelevant difficulties that arise when the objective and subjective dimensions of moral evaluation come apart, let me also stipulate that Gina is right about this.

<sup>16</sup> See, e.g., Hills (2009) and McGrath (2011). For criticism of this distinction, see Sliwa (2012). I am very sympathetic to Sliwa's criticism, as I am to her argument generally. But it is not necessary to fight that battle here.



her reliability.<sup>17</sup> While her judgement is indeed correct, it remains obscure to me why. I shall refer to this final version of the case as *Country cousin* (\*).

So I defer to Gina in the matter of what generosity requires. What is wrong with that? In general terms, a very common idea is that something morally important is lost when we defer to others. It somehow belongs to the ideal of an adult moral agent, we might say, that she works the answers to moral questions out for herself. Living up to this ideal therefore requires more of me than mere acquisition of a valid warrant to affirm some answer to my moral question [even, the correct answer], which I presumably do acquire in *Country cousin* (\*). Taking a page from Hills (2009), a central example of what I lose (or rather, fail to gain) is moral *understanding* of why generosity permits me to take the friend or floor options.

Somewhat surprisingly perhaps, I fully agree that, in some sense, it is better not to defer morally to others. Alternatively, I agree that some valuable things are lost or forgone when we defer, and even that these plausibly include moral understanding. However, I deny that this entails any good objection to moral deference. The basic point is a structural one. But we should warm up to it by noticing that the relevant structure can be framed without appealing to any especially moral values, let alone fancy theoretical ones.

Situations in which one moral agent defers to another are situations in which the first agent falls short of our ideal of moral agency. The valuable things that stand to be lost here can be analyzed as dimensions of this ideal, i.e. as respects in which someone's moral agency may prove sub-optimal. Moral understanding is one such dimension; and other contributors to the 'explain the stain' game, if I may call it that, propose various additional candidates. Yet the mundane advantages of *self-reliance* also fit the minimal bill perfectly well. That is to say, moral agents who can answer moral questions for themselves will have access to these answers even in situations where no other reliable advisor is available, whereas the same cannot be said of agents who have to rely on others to answer moral questions for them. This suffices to yield a fairly straightforward sense in which deference is sub-optimal or in which it is better not to have to defer.

However, from the mere fact that moral deference is sub-optimal, it simply does not follow that deference by some agent is at all objectionable or somehow stained or to be avoided. This basic point is demonstrable without regard to the particular

---

<sup>17</sup> We nestle up against a separate issue here, concerning how articulate (intellectually or philosophically) a reliable moral judge has to be. I deny that Gina's reliability in judgment entails that her ability to articulate the reasons for her judgments about generosity is equally reliable. But we need not engage that dispute here. Those who disagree may imagine either that (while she was capable of explaining her reasons to me) Gina was otherwise prevented from doing so or else that she did explain her reasons, but I still failed to grasp them.

dimension in which the relevant agent's moral agency proves to be sub-optimal. It will nevertheless be dialectically useful to frame the argument specifically in relation to moral understanding (rather than self-reliance, e.g.). To sharpen the point, let us also assume that moral understanding is much more important than self-reliance (as Hills 2009 seems to hold) or than any of the other candidate dimensions.

So consider moral understanding. Its ability to license objections to moral deference entirely depends on whether understanding is actually on offer in the agent's practical situation. Suppose, e.g., that I *cannot* understand why generosity does not require me to take the footing option in *Country cousin* (\*). In that case, the valuable thing that 'stands to be lost' by my deferring to Gina is *already* lost. There is nothing I can do to change that.<sup>18</sup> But then no objection to my deferring can arise (not from moral understanding, anyhow).<sup>19</sup>

Notice, crucially, that this conclusion does not result from any trade-off. More specifically, it is not the case that while moral understanding gives me *some* reason not to defer to Gina (cf. Crisp 2014, pp. 132 and 134), this reason is defeated by my reason to learn what generosity requires [which, in *Country cousin* (\*), favours deferring to Gina]. For in the present version of the case, my practical options are only to defer to Gina or to guess at the right answer or to do nothing.<sup>20</sup> None of these options will gain me any moral understanding. While it may be obvious that moral understanding therefore gives me no reason to guess or to defer, it is important to see that it does not give me any reason to do nothing either. In comparison to the first two options, doing nothing amounts to a 'dog in the manger' option, as it will certainly cost me something else of some value (and for no purpose). But moral understanding gives me *no* reason to play dog in the manger. Under the circumstances, then, it gives me no reasons at all.<sup>21</sup> Rather, moral understanding is practically inert.

This further observation explains why it is literally correct that no objection to my deferring to Gina is licensed (by moral understanding). Moreover, this conclusion is preserved under substitutions of other candidate dimensions of our ideal of moral

---

<sup>18</sup> Enoch (2014, p. 247) makes a similar observation, but does not draw the conclusion to follow in the text.

<sup>19</sup> As I read her, Hills actually concedes this point (2009, pp. 122-4). But it seems to me that this gives the game away. It no longer follows, e.g., that 'we have strong reasons neither to trust moral testimony nor to defer to moral experts' (p. 98). Hills also argues that actions performed without moral understanding (of why they are right) fail to be morally worthy. However, this is a separate point. While I reject it (too), I do not have the space to engage this further claim properly here. For some criticism of this aspect of Hills' position, see Jones and Schroeter (2012) and Driver (2015).

<sup>20</sup> Some may think that I always have a fourth option, of trying to understand and failing. Those who do should suppose that I have done that already.

<sup>21</sup> Except perhaps to regret that my situation has the features it does.

agency, as long as any such candidate (self-reliance, e.g.) is not on offer in the agent's practical situation either. Hence, my conclusion—that no objection to my deferring to Gina is licensed—holds independently of taking 'moral understanding' as the particular valuable thing that stands to be lost.

Of course, in other situations, the valuable thing might actually be on offer. In other variants of *Country cousin* (\*), e.g., I might well be able to understand the basis of Gina's moral advice. We shall come to these variants presently. But we should first recognise that the variant in which I cannot understand—or, more generally, in which the valuable thing in question is *not on offer* in the agent's practical situation—is the basic and most instructive version of the case. For that is the version in which our evaluation of moral deference is focused starkly and exclusively on comparing deference to guessing or doing nothing, i.e. to the available alternative specifications of 'not understanding' (or, more generally, of 'not the valuable thing in question'). What emerges from this comparison is the realisation that moral understanding is indifferent as between these alternatives. Its indifference is what makes it practically inert under the circumstances.

Once we see this point in the basic version of the case, however, we become better able to appreciate that it extends to other versions, too, albeit more subtly. Suppose now that I might actually manage, in *Country cousin* (\*), to understand why generosity does not require me to take the footing option. To simplify, let us say that my deliberative options are (i) to defer to Gina, (ii) to guess at the right answer, and (iii) to work that answer out for myself. While moral understanding plainly favours option (iii), it does not follow that option (iii) is in fact my best option. Whether that conclusion is warranted depends on how various considerations besides moral understanding balance out, including the likelihood of actually reaching the right answer for each deliberative alternative and the importance of reaching the right answer in the case at hand.<sup>22</sup> No doubt philosophers disagree about these questions. But to sharpen the point once more, let us simply accept that, in this variant of the case, option (iii) is my best option. (In the usual course of affairs, this conclusion would have to be earned by argument. I do not mean to be conceding that it will typically be an easy conclusion to earn.)

Against that background, consider my choice between (i) deferring to Gina and (iii) working the answer out for myself. Imagine that I decide to defer to Gina. In the specified context, my decision is gratuitous and unjustified. However, even treating options (i) and (iii) as inconsistent alternatives, my decision still has two halves, the half in which I decide against (iii) and the half in which I decide in favour of (i). Their

---

<sup>22</sup> For an argument that option (iii) can never be my best option, provided that I am more likely to reach the right answer by deferring, see Enoch (2014, pp. 248-50).

independence is secured by the existence of my option (ii) to guess at the right answer. Strictly speaking, then, the objection to my decision licensed by our case stipulations only applies to its first half, where I decide *against*(iii). That is what is gratuitous and unjustified.

Of course, one might respond that, under the circumstances, ‘having decided against (iii)’ is a *presupposition* of deciding in favour of (i). In that sense, the objection remains attached to the second half of my decision, insofar as that half inevitably carries an objectionable presupposition. While this cannot be denied, it actually helps to explain why, all the same, no objection attaches specifically to my moral deference. For the objectionable presupposition attaches equally to *anything* I decide here, once I have decided against (iii) working the answer out for myself. In particular, and by way of illustration, it would attach equally to my decision (ii) to guess at the right answer.<sup>23</sup> But any objection that applies equally to moral deference and to guessing thereby fails to discriminate between them. Hence it has nothing to do with moral deference per se.

### 3. *Coda on virtue*

When I defer to Gina in *Country Cousin*(\*), the conclusion I accept on her authority is that ‘generosity does not require me to take the footing option.’ It does not follow from this, of course, that *morality* does not require me to take the footing option. As described, then, Gina’s moral advice is partial or incomplete.

In my own view, this feature of the example is not any kind of problem. On the one hand, to put it both starkly and mildly, that is how things actually are. In life as we know it, that is, very few people are reliably knowledgeable about morality across the board. One sense, therefore, in which the example is not problematic is that it is realistic. On the other hand, despite the incompleteness of Gina’s advice, it remains the case that, in accepting it in the manner described, I fully exhibit the phenomenon of interest, namely, moral *deference*. Another sense, therefore, in which the example is not problematic is that it is still relevant to our enquiry.

Now there are various ways to bridge the gap between Gina’s incomplete moral advice and a conclusion about what, all things considered, the right thing for me to do is in *Country Cousin*(\*). In the simplest case, morality’s all things considered verdict coincides with the verdict of generosity because no other moral considerations apply in the situation at hand. In that case, all that I require to complete Gina’s advice is

---

<sup>23</sup> It is worth bearing in mind here that guessing is something an agent can do all by herself. Whatever the objection to the presupposition may be, then, it has nothing to do with the agent’s relations to others.

access to the proposition ‘no other moral considerations apply here’; as it were, access to that plank enables me to bridge the gap in question.

While the necessary plank might be procured from various sources, by far the simplest source is my own judgment (i.e., the agent’s own). For present purposes, it does not matter how reliable the agent’s own moral judgment happens to be. If mine is not very reliable, then my final conclusion about what to do in *Country Cousin*(\*) may suffer correspondingly. However, that is clearly a problem I have anyway. Our present concern is to evaluate the contribution(s) made by moral deference specifically. The fact that agents can always fall back on their own judgment to bridge any gap between the advice on offer and the conclusion they seek about the right thing to do, all things considered, means that incompleteness in the moral advice offered by an exemplar of virtue does not prevent such advice from being practically useful.

In the cases of interest to us, moreover, the exemplar’s judgment will, by hypothesis, be *more reliable* than the agent’s own, at least with respect to moral considerations specific to the virtue in question (e.g., to generosity). Hence deference to an exemplar’s moral advice, however incomplete, can only serve to improve the reliability of the agent’s final conclusion about what to do, even if that conclusion remains unreliable in certain other respects. Since improved reliability was always the primary reason in favour of deferring, we thus return to our original question about deference, namely, whether it also introduces countervailing moral difficulties. That is the question we have already asked and answered. But we are now in a position to confirm that, as I said, structuring our discussion around cases in which the advice and deference at issue are incomplete, as they are in *Country Cousin*(\*), makes no difference to the status of moral deference.

Let us turn, finally, to consider my suggestion that pleading for moral deference is easier when it is confined to the province of virtue in particular, as compared to morality at large. For those who happen to be interested in virtue anyhow, this conclusion will be no more than a happy bonus. But since most of the debate about moral deference is prosecuted in more general terms, vindicating my suggestion may still add something in behalf of the defense. I shall concentrate on two main points, both of which concern respects in which it is easier for exemplars of virtue to earn the credential of ‘expert,’ i.e. someone to whom others are licensed to defer.

For present purposes, credentialing someone as an ‘expert’ is simply a matter of his or her reliability with respect to a certain class of moral judgments. As we have seen, the central basis on which one person is licensed to defer to another (with respect to a certain class of moral judgments) is that the second person is a more

reliable judge than the first (with respect to that class of judgments);<sup>24</sup> *and* that the first person knows this. By an ‘expert,’ then, I mean no more than someone who satisfies these two conditions (cf. note 9).

My first point is that *being a reliable judge* is easier with respect to the moral considerations specific to a given virtue (generosity, say, or compassion) than it is with respect to the right thing to do, all things considered.<sup>25</sup> Up to a point, this is just a boring consequence of the fact that reliability is easier to achieve with respect to a narrower class of judgments than it is with respect to a wider class (especially when the wider class encompasses the narrower one). Beyond this, however, the extra width in the scope of judgments about all things considered rightness also covers at least two sorts of judgment differing in kind, and not merely in number, from judgments about the moral considerations specific to generosity, say.<sup>26</sup>

On the one hand, this extra width extends to judgments about the various moral considerations specific to each of the *other virtues*. (This is plainly an umbrella category). On the other hand, it extends to the judgments required to resolve *conflicts* among any number of more specific moral considerations. Even if someone is a reliable judge of generosity, then, there is no particular reason to expect that the basis of her reliability—the particular abilities that constitute it—will already equip her to be reliable about either of the other sorts of judgment, let alone both of them. This is not at all to say that an exemplar of generosity could not *develop* all of the other relevant abilities. But it is undeniably easier for her not to (have to) develop them.

Of course, the distinctions invoked here between various kinds of moral judgments presuppose that someone can *be* an exemplar of generosity or compassion without having all of the other virtues. My distinctions presuppose, in other words, that there is no unity or reciprocity of the virtues. Since I have argued for that conclusion elsewhere (Sreenivasan 2009), that is fine with me. But it may be helpful to recast my first point in these familiar terms. Under the unity of the virtues, someone can have the virtue of generosity only if she also has all of the other virtues—either because the various relevant abilities are not actually distinct existences (unity) or because, while they are, having the abilities relevant to

---

<sup>24</sup> Distinctly more reliable, if you prefer.

<sup>25</sup> The position I shall describe is broadly similar to the one advocated by Jones and Schroeter (2012), although the basis on which it is secured is somewhat different.

<sup>26</sup> Naturally, there may be additional moral considerations that are neither specific to any virtue nor common among all virtues. If there are, then all things considered rightness will include a *third* sort of judgment likewise differing in kind (or more than two sorts, anyhow). For what it is worth, it seems clear to me that there are such considerations, i.e. that morality includes more than virtue(s). Indeed, my argument against the unity of the virtues begins from this very fact (Sreenivasan 2009). But the present point does not depend upon it.

generosity, say, does *not count* as having a ‘virtue’ unless one also has all of the other abilities, too (reciprocity). By contrast, under disunity, someone can have the *virtue* of generosity all by itself. So construed, my first point is simply that it is (much) easier to have the ‘disunified virtue’ of generosity than it is to have its ‘unified’ counterpart; and likewise for any other individual virtue.

Eric Wiland (2015: 171-77) has recently argued that moral expertise is more plausibly or appropriately located at the level of ‘all things considered rightness’ than at the level of specific judgments of compassion or gratitude, say.<sup>27</sup> On the face of it, his claim runs precisely counter to mine. Nevertheless, this appearance is misleading in one notable sense. In another sense, it is not misleading, but does correspond to an important point that Wiland has overlooked. It will be useful to clarify these matters briefly before proceeding to my second main point.

To begin with, Wiland actually agrees that judgements about all things considered rightness are much more difficult than specific judgements of compassion or gratitude.<sup>28</sup> Indeed, that is his point of departure. For specific judgements of gratitude (say) are not only easier, he claims, but in fact ‘self-evident,’ so that only a ‘moral idiot’ would require help to make them properly. That is what is supposed to make deference to experts at this level *inappropriate*; and Wiland’s fundamental observation is that this ‘problem’ does not arise at the level of all things considered rightness. Nor do I think that the greater difficulty attending judgements of all things considered rightness makes claims of expertise or deference to experts at this level in any way ‘inappropriate.’ Rather, my contention—to foreshadow the lesson to follow below—is only that it makes experts about all things considered rightness very scarce (and hence, less available to help the rest of us). I also agree that some specific judgements of gratitude (or generosity or whatever) are reasonably regarded as self-evident; and consequently, that few normal adults will be in position to require advice about them (at least, in ‘our’ culture).

What Wiland overlooks, however, is that *not all* specific judgments of this or that virtue are self-evident.<sup>29</sup> Hence, it is not necessary to ascend to the level of all things considered rightness to find moral terrain on which normal adults may, with perfect propriety, find themselves in need of moral advice. More specifically, within the class of judgments specific to a given virtue, we can distinguish ‘paradigmatic’ judgments

---

<sup>27</sup> Wiland himself formulates the (lower) level of specific judgments in the language of Rossian ‘prima facie duties,’ rather than in the language of virtue. Yet, as he recognizes explicitly (2015: 171), nothing turns here on the details of the formulations.

<sup>28</sup> Following Ross, he focuses on the dimension of *conflicts* among prima facie duties.

<sup>29</sup> The same goes for prima facie duties. Wiland’s twin contrasts—‘prima facie duties’ versus ‘all things considered rightness’ and ‘self-evident’ versus ‘not self-evident’—therefore fail to align neatly in the way that his argument requires.

of that virtue from ‘non-paradigmatic’ judgments. *Paradigmatic* (or ‘stereotypical’) judgments of some virtue are judgments that anyone who is acquainted with that particular virtue concept can make; and they are typically also the judgments by reference to which the concept is taught. ‘Helping an old lady to cross the street,’ for example, is a paradigmatic act of kindness. Anyone who is familiar with the concept of kindness knows that it is ‘kind’ to help [suitably situated] old ladies to cross the street. Wiland is certainly right that no such person requires moral advice to make *that* judgement.<sup>30</sup>

By contrast, *non-paradigmatic* judgments of kindness are simply moral judgments that are plausibly specific to kindness and yet not paradigmatic of it. It may be controversial exactly where (and how) to draw the line between ‘non-paradigmatic judgments of kindness’ and moral judgments that not judgments ‘of kindness’ at all. But I take it, in any case, that paradigmatic judgments of kindness clearly do not exhaust the class of judgments specific to kindness. So it is not controversial that there are some non-paradigmatic judgments of kindness. Along with their counterparts for other individual virtues, these judgments satisfy both Wiland’s criterion that someone may stand in need of moral advice about them without qualifying himself as a ‘moral idiot’; and mine that reliability about them is easier to achieve than reliability about judgments of all things considered rightness.

Recall that an ‘expert’ about kindness, in the present sense, is someone who is both *reliable* about making judgments of kindness and *known to be reliable* on this score.<sup>31</sup> My second main point is that this second condition is also easier to satisfy within the province of virtue. That is to say, not only is reliability in judgment itself easier to achieve, but so too is a *reputation* for reliability—or, more precisely, the accessibility of the evidence for someone’s reliability—easier to secure with respect to individual virtues. As I said earlier (see note 14), there is some question about how those of us in need of moral advice are supposed to go about identifying moral experts in the first place. Indeed, various skeptics about moral deference seem to feel that this represents some kind of insuperable stumbling block.

In general terms, the challenge is to find an evidence basis for someone’s judgmental reliability that satisfies all of the following desiderata: First, it really is

---

<sup>30</sup> It does not follow that, when co-located with an old lady and a street, everyone familiar with the concept of kindness always knows whether to help her across or not (not even when no other moral considerations apply). For there is still a significant difference between making a moral judgment in a stylized ‘textbook’ context and making the ‘same’ judgment out in the world. See, e.g., the persuasive account in Blum (1994). Accordingly, even within the core of paradigmatic judgments for a given virtue, there remains room for differential reliability in practice between normal adults, and thus room for moral deference.

<sup>31</sup> For simplicity, I elide the comparative dimension of expertise here.



evidence of reliability (so that trust on its basis is neither blind nor arbitrary); second, one does not have to be an ‘expert’ oneself to employ it; and third, it is not itself controversial. The argument (or rather, claim) in the literature concerning possible asymmetries between the moral and non-moral domains in relation to expertise has been that, in the moral domain, these desiderata either cannot be satisfied or are very difficult to satisfy jointly.

At least within the province of virtue, I claim that a straightforward basis for identifying moral experts is made available by the distinction between paradigmatic and non-paradigmatic judgements.<sup>32</sup> Consider kindness again. My suggestion is that experts in kindness can be identified on the basis of their (greater) reliability in performing paradigmatically kind acts. These acts can be identified as acts of kindness by *anyone* who is familiar with the concept. Hence, access to this (potential) evidence does not depend upon one’s being an expert oneself. Moreover, the open accessibility of paradigmatically kind acts still leaves room for differentially reliable performance between normal adults (cf. note 30). The rest of us are therefore in position, more specifically, to recognize others—if or when we encounter some—who are *more* reliable about acting kindly than we are (i.e., others who have comparative expertise or expertise, proper). Finally, the status of paradigmatically kind acts as ‘kind’ is not controversial in any relevant sense.<sup>33</sup> Accordingly, paradigmatically kind acts or judgments satisfy all three desiderata for an evidence basis on which to identify reliable judges of kindness.

The scope of the licence to defer to experts in kindness one acquires on this basis plainly includes paradigmatically kind judgments in all their manifold variety in practice. It would also be reasonable, however, to treat the scope of this licence as extending further, to non-paradigmatic judgments of kindness as well. Unlike judgments of *other* virtues (such as generosity or gratitude), non-paradigmatic judgments of kindness respond to the *same kind* of moral considerations as paradigmatic judgments of kindness. The particular abilities that constitute a reliable agent’s reliability with respect to the latter judgments are therefore, I submit, reasonably seen as enabling her to be reliable with respect to the former

---

<sup>32</sup> Whether there are (also) good bases for identifying experts in the wider moral domain I simply leave open. Thus, by ‘*easier* to identify experts in particular virtues,’ I mean that we already have a straightforward basis for doing so, as argued in the text, whereas it remains obscure what basis there might be for identifying experts within morality more generally.

<sup>33</sup> Their status as kind is not controversial within the community of those who endorse the operative concept of kindness, which community includes both experts and those in need of moral advice. Of course, as with any moral notion, ‘kindness’ may be controversial in other, deeper senses (e.g., to moral skeptics and nihilists or between different moral cultures).

judgementstoo.<sup>34</sup> In addition, we should notice that when I defer to an expert's non-paradigmatic judgments of kindness, what I am deferring to her about is, in effect, the question of how best to generalize the application of certain moral considerations we both recognize from central cases on which we agree to neighbouring cases where I am unclear. Thus, while it is still clearly *moral* deference, my deference here has a less radical flavour insofar as it is confined within the bounds of a shared evaluative perspective.

Let us take stock. My plea for moral deference has two parts. Its main part, recapitulated below, consists in explaining why there is nothing objectionable about deference, notwithstanding the consensus in the literature. This explanation holds equally in all provinces of morality. But its usefulness presupposes that those of us in need of moral advice stand a chance of actually finding someone to whom we are licensed to defer (about the matter at hand). The second part of my plea concerns the satisfaction of this presupposition; and that is the part it is easier to deliver within the province of virtue. A licence to defer to someone depends both on that person (who must be a reliable judge) and on the agent himself (who must have a secure basis on which to identify another's reliability). In this coda, I have argued that each of these conditions is more easily established with respect to individual (disunified) virtues. Experts in these virtues are consequently more available to the rest of us than experts in all things considered rightness.

#### 4. *Conclusion*

I have sought to defend moral deference against a prevalent skepticism. My main argument has been that the mere fact that the circumstances in which deference occurs are typically sub-optimal does not ground any good objection to the agent's deference itself. Someone who defers to another person may not fully achieve our ideal of an adult moral agent. Fair enough.<sup>35</sup> Still, this fact does not result from the agent's deference nor does it prevent her deference from being a blemish-free response

---

<sup>34</sup> The inference from 'reliable about paradigmatic judgments' to 'reliable about non-paradigmatic judgments,' all within the same virtue, is clearly *more reasonable* than the parallel inference from 'reliable about one [disunified] virtue' to 'reliable about another [disunified] virtue.' However, my claim is that the first inference is also reasonable tout court. I argue for this claim elsewhere (Sreenivasan, forthcoming).

<sup>35</sup> While I suspect that skeptics also greatly over-value the full achievement of their ideal of agency, that is another matter.

to the agent's circumstances as they actually are. Often, deference will also be her best response.\*

## *References*

- Blum, L. (1994) *Moral perception and particularity*. Cambridge: Cambridge University Press.
- Crisp, R. (2014) "Moral testimony pessimism: A defence," *Proceedings of the Aristotelian Society, Supplementary Volume* 88: 129-43.
- Driver, J. (2006) "Autonomy and the asymmetry problem for moral expertise," *Philosophical Studies* 128: 619-44.
- (2015) "Virtue and moral deference," *Ethics & Politics* (this issue).
- Enoch, D. (2014) "A defense of moral deference," *Journal of Philosophy* 111: 229-58.
- Hills, A. (2009) "Moral testimony and moral epistemology," *Ethics* 120: 94-127.
- Hopkins, R. (2007) "What is wrong with moral testimony?," *Philosophy and Phenomenological Research* 74: 611-34.
- Howell, R. (2014) "Google morals, virtue, and the asymmetry of deference," *Noûs* 48: 389-415.
- Jones, K. (1999) "Second-hand moral knowledge," *Journal of Philosophy* 96: 55-78.
- Jones, K. and Schroeter, F. (2012) "Moral expertise," *Analyse&Kritik* 2: 217-30.
- McGrath, S. (2011) "Skepticism about moral expertise as a puzzle for moral realism," *Journal of Philosophy* 107: 111-37.
- Sliwa, P. (2012) "In defense of moral testimony," *Philosophical Studies* 158: 175-95.
- Sreenivasan, G. (2009) "Disunity of virtue," *Journal of Ethics* 13: 195-212.
- (forthcoming) "Acts, agents, and the definition of virtue." In M. Timmons (ed.) *Oxford Studies in Normative Ethics*, volume 7. Oxford: Oxford University Press.
- Taylor, C. (1989) *Sources of the Self*. Cambridge, MA: Harvard University Press.
- Wiland, E. (2015) "Rossian deontology and the possibility of moral expertise." In M. Timmons (ed.) *Oxford Studies in Normative Ethics*, volume 4. Oxford: Oxford University Press, pp. 159-78.

---

\* Earlier versions of this paper were presented, sometimes along with other material, at the annual meeting of the Canadian Philosophical Association in Victoria; at a philosophy colloquium at the CUNY Graduate Center; and at the *Varieties of Virtue Ethics* conference, held in Oxford by the Jubilee Centre for Character and Virtues. Thanks to the audiences on all three occasions for helpful discussion. For comments on material specific to this paper, I am especially grateful to Sarah Buss, Julia Driver, David Enoch, and Eric Wiland.