



THE  
JUBILEE CENTRE  
FOR CHARACTER & VIRTUES

UNIVERSITY OF  
BIRMINGHAM

## **“A plea for moral deference”**

**Gopal Sreenivasan**

This is an unpublished conference paper for the 3<sup>rd</sup> Annual Jubilee Centre for Character and Virtues conference at Oriel College, Oxford University, Thursday 8<sup>th</sup> – Saturday 10<sup>th</sup> January 2015.

These papers are works in progress and should not be cited without author's prior permission.



“A plea for moral deference”

**Gopal Sreenivasan**

Crown Professor of Ethics  
Duke University

[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 scepticism. I hope to contribute to rehabilitating the notion for some role in moral theory.

## A plea for moral deference

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

---

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

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

---

or ‘fishy’ (Enoch 2014).

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

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 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 (i.e., easier than the case we shall already have successfully observed).

---

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

## §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 democratisation 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 reflection, moreover, this particular idea is indeed 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

---

<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.

to moral expertise appears to be inconsistent with the objectivity of morality.<sup>6</sup> Whether or not that counts as a 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.

---

<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.

<sup>7</sup> Here as elsewhere, we may be reminded of the familiar symmetry between modus ponens and modus tolens. 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.



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.

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 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

---

<sup>8</sup> We can refine the relevant conditions 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.

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 recognise 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 quality and differential quantity (of experience). 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

---

<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

his elders still), i.e. our raw fact. To resist this conclusion, one would have to deny that this differential typically does persist, 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 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

---

learning, in the next paragraph, to cover just these points.

<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.

something yet more specific. With this case, it would be natural, e.g., to wonder whether *generosity* 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 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,

---

<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.

<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 italicised concession. Nevertheless, a profitable question does remain to be discussed here, and I shall return to it (too) in the coda on virtue.

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 presents an example precisely of 'deference,' so understood, let us specify that I accept Gina's judgement here simply on the basis of her reliability.<sup>17</sup>

---

<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.

<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 judgement entails that her ability to articulate the reasons for her judgements about generosity is equally

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 foe 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, and also filled in, 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

---

reliable. But we need not engage that dispute here. Those who disagree may imagine either that (while she was capable of doing so) Gina was otherwise prevented from explaining her reasons to me or else that she did explain her reasons, but I still failed to grasp them.

lost here can be analysed as dimensions of this ideal, i.e. as respects in which someone's moral agency may prove sub-optimal. Moral understanding is one of those dimensions; and other contributors to the 'explain the stain' game, if I may call it that, have proposed 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 simple sense in which deference is sub-optimal or in which it is better not to have to defer.

Since our ideal of moral agency is plausibly multi-dimensional, the various proposed candidates need not be regarded as competitors: each of them may capture a (distinct) sense in which deference is sub-optimal. Still, it is consistent with this possibility that some dimensions of the ideal are (much) more valuable or important than others. Thus, to sharpen the main point, let me simply concede that moral understanding is much more important than self-reliance (as Hills 2009 seems to hold) or even than any of the other candidates.

Consider moral understanding then. Its ability to license objections to moral deference entirely depends on whether it 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. But then no objection to my deferring arises (not from moral understanding,

anyhow).<sup>18</sup>

Notice, crucially, that this conclusion does *not* result from any trade-off. It is not the case, more specifically, that while moral understanding gives me *some* reason not to defer to Gina, 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 do nothing.<sup>19</sup> Neither option will gain me any moral understanding. Doing nothing therefore amounts to a ‘dog in the manger’ option, since it will cost me something else of some value. But moral understanding gives me *no* reason to play dog in the manger. As a result, it gives me no reasons at all;<sup>20</sup> rather, moral understanding is practically inert here.

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 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

---

<sup>18</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).

<sup>19</sup> Some may think I have a third option, of trying to understand and failing. Those who do should suppose that I have done that already.

<sup>20</sup> Except perhaps to regret that my situation has the features it does.



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 the question, to defer or not to defer.

Once we see the point in the basic version of the case, however, we gain the position to appreciate that its core 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 options concerning how to proceed deliberately are (i) defer to Gina, (ii) do nothing, and (iii) work the 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 with different options and the importance of reaching the right answer in the case at hand. 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.

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 independence is secured by the existence of my option (ii) to do nothing. Strictly speaking, then, the stipulatively licensed objection to my decision only applies to its first half, where I decide *against (iii)*. That is what is gratuitous and unjustified. Even here, in other words, the objection does not apply to my moral deference per se, for option (i) is no more liable to this objection than option (ii) would be. Thus, if we held the first half of my decision fixed, and evaluated my decision to defer under that assumption, no objection to my deferring would remain (indeed, deferring would be highly preferable to my doing nothing).

In summary, I am happy to concede that it is better not to have to defer to others on moral questions. But it does not follow that moral deference is at all objectionable or to be avoided.

### §3. *Coda on virtue*

This section remains incomplete, alas.

## References

- Driver, J. (2006) "Autonomy and the asymmetry problem for moral expertise," *Philosophical Studies* 128: 619-44.
- Enoch, D. (2014) "A defense of moral deference," *Journal of Philosophy*
- 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. (2009) "The puzzle of pure moral deference," *Philosophical Perspectives* 23: 321-44.
- (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.
- Taylor, C. (1989) *Sources of the Self*. Cambridge, MA: Harvard University Press.