

*Google Morals, Virtue, and the Asymmetry
of Deference*

ROBERT J. HOWELL

Southern Methodist University

Although I don't know the streets and subways of New York very well, I get around with little difficulty. And if I'm suddenly struck by a desire for tofu bahn mi, no sweat. Like many of you I have an application on my phone, Google Maps, which gives me directions and directs me to restaurants at the push of a few buttons. I am therefore somewhat dependent on my phone, but so what: it's a perfectly reasonable way to get around.

Suppose those wizards at Google come out with a new app: Google Morals. No longer will we find ourselves lost in the moral metropolis. When faced with a moral quandry or deep ethical question we can type a query and the answer comes forthwith. Next time I am weighing the value of a tasty steak against the disvalue of animal suffering, I'll know what to do. Never again will I be paralyzed by the prospect of pushing that fat man onto the trolley tracks to prevent five innocents from being killed. I'll just Google it. Again I find myself dependent on Google for my beliefs, but in this case it seems to many, myself included, that this is not a good way to go. There seems to be something wrong with using Google Morals. But what is it?

The issue, of course, is whether or not there is something particularly wrong with forming beliefs based on deference in moral matters. Is there an asymmetry of deference between the moral and the non-moral? If there is, what is the nature of the asymmetry? Is it epistemic, moral, or something else?¹

I have the sense that there is an asymmetry, and that this asymmetry remains even if moral judgments can be true and objective, in the sense that they are not made true by the state of mind of the judge. This asymmetry persists even if there are quite a few moral experts in the sense that there are people who know the objective truth about a broad range of moral matters.²

I am on the hunt for an explanation of the asymmetry, therefore, that is consistent with these assumptions.

Degrees of Deference

Deference involves forming and sustaining a belief in *p* based upon some other act of believing *p*. Two features of this definition are worth noting. First, deference involves what *sustains* belief. No doubt it can happen that we ask for advice, and in getting that advice we see the reason behind it. We also might wind up integrating the advice into our own reasons. In such cases, the word of others might prompt belief, but it does not in the end sustain it, so it is not what I am calling deference. Second, deference usually involves forming and sustaining a belief based on the fact that someone *else* believes, but that is inessential—one could also form and sustain a belief that *p* based on the fact that either one did believe or will believe that *p*. One could have had brain damage and lose one's moral beliefs, but find out one used to believe that *p*. One could then defer to that belief and judge that *p*. Or one could learn somehow that one would in the future judge that *p* and therefore judge that *p*. The asymmetry of deference seems to exist here as well. If I have amnesia, and I discover that I used to believe that my wife's name was Lanie, I should probably defer to my past self. If, on the other hand, I have "moral amnesia" and suddenly don't recall my old values, but I discover I used to believe that abortion is wrong, it seems problematic for me to adopt that belief. An account of the asymmetry of deference should account for what is problematic about self-deference as well as deference to others.

What does it mean to say that moral deference is problematic? There are several issues here. Does it mean that something is wrong with the act of deference, or that something is wrong with the agent who defers? Does it mean that deference is wrong, that it is *pro tanto* wrong, or simply that it is somehow sub-optimal?

It simply cannot be the case that the act of deferring is always wrong. Normal agents should certainly defer in extreme cases, when circumstances are exigent and the costs of inaction are dire. But other agents should defer even in more common circumstances. A sociopath should often defer to others. But we needn't go to sociopaths to find agents who ought to defer pretty much across the board. Young children should and indeed must defer if they are to become moral agents at all. In such cases, deference isn't even a *pro tanto* bad. Any diagnosis of the asymmetry of moral deference must accommodate these facts.

Still, even if one must admit that in some cases there is nothing bad about deference, it might seem that for normal agents in normal circumstances, the act of deference is wrong when it comes to moral matters. It is far from clear that this is the case, and the asymmetry of deference need not have this

implication. If one is in a position to learn a moral fact by deferring, and one cannot come to know that fact non-deferentially without substantial cost, it might well be that almost always one should do so. Otherwise, one is apt to do something morally impermissible. There is still room for an asymmetry of deference, however. An asymmetry is preserved just so long as deference is somehow sub-optimal to the extent that it would have been better if the knowledge had been attained without deference.³ It could be that this value of non-deference is almost always outweighed by the risk of performing an impermissible act, but that value remains.

Is the act of deferring sub-optimal, or is the agent sub-optimal? Given that in many cases where one needs to defer one should, it is difficult to find fault in the act.⁴ For reasons that will become clear, I'm inclined to think that deference shows that something is sub-optimal in the agent. This still leaves open an important question. Is the need to defer indicative of something that was already problematic in the agent, or does the deference lead to a problem in the agent? In fact, I will maintain that both can be the case and each will play a role in the explanation of the asymmetry.⁵ Importantly, though, the agent is not sub-optimal simply because she is initially ignorant and needs to defer. That much is trivial, and would not mark any asymmetry between the moral and non-moral cases.

There are interesting questions as well about the scope of the problem. Is deference problematic only in the case of morality, or is it similarly problematic in aesthetics and other normative domains? What is fishy about moral deference seems similarly fishy in aesthetic deference, and perhaps in other normative domains as well.⁶ Deference in the moral domain might be particularly troublesome because a subject's moral beliefs are more important than her aesthetic beliefs, but deference in both domains is problematic for the same reasons. A distinction might remove the appearance to the contrary. One could do at least two things in deferring to another: one could defer in action, or one could defer in belief. If I have a shaky sense of style, I might ask my wife which shirt looks best on me, and I might invariably take her advice by wearing the shirt she recommends. This might involve proceeding as though her judgment is true without forming a belief about whether the shirt really looks best. Call this *active deference*. Many cases of deference, I think, fall into this category, and it's plausible that nothing is wrong with active deference in the aesthetic case while there is in the moral case. There is also *doxastic deference*, however, in which the deferring agent forms a belief. It is doxastic deference that I am concerned with, and here the aesthetic and moral cases both seem problematic. In doxastic deference the agent takes an assertoric attitude towards the proposition, believing it. While there doesn't seem to be anything wrong with my taking my wife's advice, it seems strange for me to assert that, say, plaid looks great on plaid, if I don't *see* it.

There are also many different varieties of moral deference, and they do not all seem equally problematic. Some cases involve deference in a moral matter,

but the deference is really to another's expertise regarding non-normative facts. One might defer to someone about the impermissibility of eating meat, for example, because that person had either worked out the consequences of vegetarianism—which involves knowing all sorts of things including how abstinence affects markets—or because that person had superior scientific knowledge about the capacity of certain animals to suffer. These cases involve taking a moral stance because of deference, but the deference really concerns the non-moral expertise. Let's call such deference *empirical moral deference*.

Perhaps there is something wrong with some cases of empirical moral deference. If, for example, the stakes are terribly high morally, it might be irresponsible to take someone's word even for the empirical facts. Nevertheless, this seems quite different from purer forms of moral deference.⁷

There is important diversity even among these "purer" forms of deference. There are different reasons for going to your priest for moral guidance. You can seek advice knowing that you share basic moral values, asking what to do in light of those values. You might do this because of his greater practice and sophistication at weighing the values in your system, at seeing how they interact and offset one another, and at applying those values to cases. So, you might ask "Should I lie to my parents when they ask whether or not I think they were good parents?" On the one hand you believe you shouldn't lie, but on the other you believe you must honor your father and mother. Here your priest is acting as a sort of moral calculator, tallying up the points for and against the action using your own values. We might call such a case *derivative moral deference*, because it is moral deference but it is ultimately grounded in your own values.

Derivative moral deference itself comes in flavors.⁸ In the example above, it was clear which of your values were at play, you just didn't know which one prevailed. But you might know the priest is using your values, without knowing which ones or how the moral calculations are being made. This is a sort of *derivative moral deference with ignorance*. And, of course, there is a further case where you don't know your advisor is using your values, but in fact he is using them. The deference is still derivative, in a sense, because you are deferring to a moral verdict that is based in your values, you simply don't realize that fact.⁹

There are also cases of what we might call *foundational moral deference*. One might not be going to a priest to see how one's antecedently held values apply to the case at hand. One might go to the priest to find out what is right or wrong, so that in deferring to the priest one embraces a value one did not previously have, and that doesn't derive from values one has.

One could no doubt go further and develop a Linnean taxonomy of deference, but the main point should be clear. There are many difference types of deference, and it is not obvious that all of them are equally problematic. In fact, the cases seem increasingly problematic in the order I presented them: empirical deference seems unproblematic, derivative forms of deference

are somewhat troubling (more so as the deferring agent is ignorant about the source of the moral judgment) and the most troubling cases involve foundational deference.

The varieties of deference open up the possibility that there might be different reasons why different forms of deference are problematic. A unified account would be nice, but it's not obvious from the outset that such an account will be available.

Finally, let me be clear about what constitutes an asymmetry of deference. It does not involve saying that there is never anything wrong with non-moral deference while there is always something wrong with moral deference. Nor, I think, should the claim be that an agent who defers in moral cases is worse off all things considered than she would be if she refrained from deferring. The claim is rather that there is something wrong with moral deference that is not in general wrong with non-moral deference, and that this fact makes moral deference worse than non-moral deference, all things being equal.

Epistemic Accounts of the Asymmetry of Deference

It is extremely tempting to explain the asymmetry of deference in terms of an epistemic asymmetry stemming from idiosyncrasies about the way moral knowledge can be gained or passed on. Assuming moral truths can be known, three different sorts of epistemic asymmetries suggest themselves:

- a. One cannot come to know moral truths by deference.
- b. It is particularly difficult to come to know moral truths by deference.
- c. One can know moral truths by deference just as easily as one can other truths, but there is a further epistemic good that one cannot get by deference.

In the current literature, there don't seem to be many defenders of a, that one cannot come to know moral truths by deference. If moral truths are objective and knowable, it is difficult to see why one couldn't come to be in a secure epistemic status with respect to those truths using the testimony of other agents. There are, however, quite a few defenders of b and c. I think b accounts are problematic and that c accounts are still problematic but less so. Both accounts, though, fall short of giving us the whole story.

Challenges for Deferential Moral Knowledge (B-Accounts)

Even if moral knowledge is attainable by testimony or deference, it might come less easily than deferential knowledge of non-moral matters. One possibility is that in the case of moral matters it is more difficult to know that one's testifier is trustworthy.¹⁰ There could be at least two reasons for lack of trust: the testifier might not know what he is talking about, or the testifier might, because of partiality or a desire to manipulate, lie about what he

believes. In the case of non-moral matters, it might be said, there are many external markers of reliability and there are fewer reasons to suspect the motives of the testifier. Given these facts, moral testimony brings risks that do not threaten other sorts of testimony.¹¹

This account has several problems, but we can start by noting how strange it is to apply this sort of account to cases of self-deference. Worries about self-manipulation are far fetched, and it would be odd to refrain from deferring because one didn't have evidence of one's own reliability.¹² In self-deference, any reason to suspect the competence or reliability of the person one was deferring to would be a reason to suspect one's ability to get to the truth oneself. Perhaps there are cases in which one can imagine having particular doubts about one's past (or future) reliability, but deference remains problematic even when one doesn't. The epistemic account thus can't explain what is wrong with self-deference.

This sort of account seems to have two further problems. On the one hand it's not clear that moral matters really are unique in these respects. Even if they are, however, this account depends on an implausibly demanding picture of knowledge transmission. At the very least, it depends on a picture that many find unattractive in non-moral matters, and if it is somehow more appealing when it comes to moral matters we have really just pushed the question of asymmetry back: why does one model seem to hold for moral matters but not for others?

First, let's examine the alleged asymmetries. There are two distinct reasons, on this account, why it might be more difficult to tell that a testifier is trustworthy. First, there is the issue that moral matters involve high stakes. This could matter either because high stakes raise the standards for knowledge or that the stakes introduce greater possibilities of partiality and motive to manipulate.¹³ Whatever the role of stakes, there is no asymmetry with the non-moral domain. Many moral matters are low stakes and many non-moral matters are high stakes. There is a low stakes moral question as to whether it is morally wrong to burp at a table, and there is the high-stakes non-moral question of the strength of the riveting in an airplane's wings. For every moral question of a certain level of importance, it is likely that there is a corresponding non-moral question that is just as important. (If the normative supervenes on the non-normative, for every normative issue there would be corresponding non-normative matters. The same holds if a moral imperative requires that some non-normative condition obtain.) These same considerations, I think, provide reasons to suspect impartiality and beneficence in the relevant non-moral cases. One has only to look at hot button topics like climate change and evolution to see how motives might taint objective investigation and reportage in empirical cases that are perceived to be morally relevant. Yet we do, and we should, defer to experts about climate change and evolution. Finally, the issue of manipulation is a bit of a red herring. Suppose instead of deferring due to an agent's testimony, one was

able to determine what the agent really thought—perhaps using some great fMRI of the future. Accepting this other agent’s belief as one’s own remains problematic despite the lack of opportunity for deception and manipulation.

The second issue is whether it is particularly difficult to tell if a moral testifier is an expert. There are interesting reasons why this might be the case. Establishing an agent’s reliability often involves getting a “track record” for that agent by establishing that the agent has succeeded in the past. But to establish that the agent has succeeded, it seems one must already be able to detect the relevant moral truths.¹⁴ So, either one cannot establish reliability or one has no need to defer.

Again, the asymmetry here is not as plausible as it first appears. There are cases when a moral expert is easier to find without being redundant and there are cases where markers of non-moral expertise are difficult to come by. First, it is not necessarily the case that one has no need to defer if one is able to ascertain a track record for an expert. One could be a good moral judge with respect to a certain range of values, develop a track record based on that knowledge, and then defer on a matter outside of that range. Additionally, one can come to know someone is reliable without oneself establishing their track record. For one thing, we can accept testimony about their reliability. We trust scientists all the time based on such testimony. In these cases, perhaps, we know that they have prestigious positions at research institutions, or that they publish in reputable journals. But usually our knowledge that those research institutions or journals have a good track record relies on testimony.

In addition, track records are hard to come by in many non-moral cases. In the highest levels of science there is often a very small group of experts in a position to evaluate the track records of each other. If I could evaluate their track records, I probably wouldn’t need to defer to them. Yet I do defer to them, and there is nothing wrong with this.

I have been giving reasons to doubt that there are particular difficulties in establishing the trustworthiness of ostensible moral experts. This only matters on the questionable assumption that we have to establish their reliability. In general, we do not require that a source of knowledge be known to be reliable in the first place. If we required this in general, we would have a short path to skepticism. If knowing *p* by source *S* required first knowing that *S* was reliable, then that knowledge, that *S* was reliable, would require that we know of its source *S'* that it was reliable. And so on. This generates a vicious regress, making any knowledge impossible. If there is knowledge, some sources must be basic, in that they can produce knowledge even if we do not know that they are knowledge producers.

Of course it doesn’t follow that testimony is a basic source of knowledge. But even if it is not really basic, we often don’t require independent and prior knowledge of the reliability of our testimonial sources. This is in part because many think that we do typically have background information that

affirms the reliability of our sources.¹⁵ If we did require knowledge of the reliability of our sources, we would surely know much less than we take ourselves to know. We often form beliefs based on the say-so of strangers. If I stop and ask for directions, I might avoid asking someone who looks patently unreliable, but I will often believe someone whom I don't know to be reliable. Absent obvious defeaters, people typically take other people's words for things and we don't tend to think there is anything wrong with that. We do feel some unease about people's accepting moral testimony. Perhaps we become more internalist when we are talking about moral matters. But this only relocates the asymmetry: why is it that we require knowledge that our source is reliable in the moral realm but not in the non-moral realm?¹⁶

Perhaps the advocate of this epistemic strategy could argue that reliable testimony is a source of justified beliefs, *assuming there are no defeaters*. In the moral case, however, there are defeaters. We know, after all, that there is extensive disagreement about such matters and that there is no systematic way to resolve such disagreement. So it could be that in moral cases, this provides a sort of undermining defeater for testimony as a source of knowledge.

This is a pretty plausible story, but it has several shortcomings. First, there are many moral matters on which there is little to no disagreement. Almost no one finds it acceptable to torture people for fun. Kantians and consequentialists might disagree about the reasons why, and there might be tricky cases in the neighborhood, but within a certain domain there is widespread agreement, yet deference in these matters is if anything worse than in the controversial matters. Of course there will remain some who will disagree even with the most foundational moral claims, and convincing them will be difficult. But in the non-moral domain there are likely to be holdouts on obvious cases as well. We don't take this to undermine knowledge transmission. Second, the undermining effects of disagreement are most plausible when the subject *knows* about the disagreement and has no independent evidence helping her choose an expert. But take a subject who does not know that there is widespread disagreement and who happens to have a moral expert at hand to whom she defers.¹⁷ This subject is in a position to know by deference, even on controversial matters but her situation seems no less suspect; the problem with deference remains.¹⁸ Finally, if the fact of disagreement is devastating to knowledge, not only might it become impossible to get by testimony alone (which is more than the b-theorist desired) but it might make it impossible to have moral experts at all. If moral testimony is particularly hard to get because of something that characterizes the moral domain in general, then to the degree that moral testimony is undermined, sources such as moral intuition are similarly undermined. But the problem of asymmetry was raised on the assumption that there were moral experts and moral skepticism was false. If the defense of the epistemic strategy tilts us towards skepticism we have a reason to think it is not the diagnosis we are looking for.

In general, belief formed on the basis of moral testimony or deference can satisfy the standards we hold for knowledge in other domains. But even if it doesn't and this explains some of our reluctance to defer, that cannot be the whole story. For on the assumptions we have made, that moral judgments are cognitive and objective, the absence of readily identifiable experts is at most a merely contingent feature of our world. But even were there a Google Morals which we knew to be reliable, there would still seem to be something wrong—or at least sub-optimal—with those who simply deferred to it. The difficulty of attaining moral knowledge can not fully account for the asymmetry of deference.

Moral Understanding

Several philosophers have claimed that moral deference provides knowledge but that such knowledge leaves something to be desired. Most compelling is the view that moral deference brings knowledge but not *understanding*. Someone who believes something by deference might know that it is true, but they do not understand *why* it is true. Furthermore, understanding why something is true is of particular importance in moral contexts. Among other things it is required for producing virtuous agents and for producing reliably good actions.¹⁹

What exactly is involved in understanding? Understanding, like knowledge, is factive: one cannot understand p if p is false. But understanding does not simply involve propositional knowledge. It involves a range of abilities. Alison Hills provides a compelling case that

If you understand why p (and q is why p), then in the right sort of circumstances, you can successfully:

- i) follow an explanation of why p given by someone else;
- ii) explain why p in your own words;
- iii) draw the conclusion that p (or that probably p) from the information that q ;
- iv) draw the conclusion that p' (or that probably p') from the information that q' (where p' and q' are similar to but not identical to p and q);
- v) given the information that p , give the right explanation, q ;
- vi) given the information that p' , give the right explanation, q' . (Hills, 102–3)

It is extremely likely, in my view, that in many cases of moral deference one will not thereby have understanding. If I simply Google “Is causing pain to non-rational animals wrong?” and I get the answer “yes!”, I will not necessarily be able to say why. My belief will not necessarily be incorporated into a system of morality that allows me to recognize the reason why the lack of rationality is irrelevant to whether or not the pain of animals is bad. After deferring if I was presented with the question “Is causing pain to animals without self-conceptions wrong” I wouldn't necessarily see the connection

between the questions. So, I wouldn't be able to defend or explain my belief to others and I wouldn't be able to apply my belief to new and similar cases. My inability to do this will probably affect my ability to be moral in the long run, and it certainly affects my ability to be a legitimate judge of others, or even a citizen who participates in a system that governs others.

A worry about this strategy is that while there is a clear difference between knowing *p* and understanding why *p*, it is not as clear that there is a difference between understanding why *p* and knowing why *p*. Hills insists that understanding involves certain know-how over and above knowledge-that, but knowing why *p* would typically involve many of these abilities as well. In general, the distinction between know-how and knowledge-that doesn't show that knowledge-that can always be had without know-how.²⁰ There might be conditions preventing knowledge-that from giving know-how—if the knower is deaf or easily distractible, for example. But these conditions are beside the point—what is important is that the understander has the cognitive where-withal to provide explanations, etc., even if she is unable to do so. It seems likely that knowledge why *p* must provide this cognitive capacity and that a lack of such capacity indicates a lack of knowledge.

Though I think this is troublesome for the understanding strategy as it stands, a modification can preserve the strategy's insight. Instead of insisting on the distinction between knowledge and understanding, the account should claim that knowing *p* (at least for normative propositions) doesn't give knowledge why *p*, and that this is especially important in moral matters. This weakens the strategy a bit, since if understanding boils down to a certain sort of propositional knowledge, then understanding can be attained through deference, *pace* Hills. But this is only a problem if someone who has understanding but has it by deference is still sub-optimal. It is not at all clear that this will be the case. We all started out our moral careers by largely deferring to our elders. It's plausible that what starts out as feeble moral knowledge becomes less so precisely by being supplemented by knowledge why, even if that knowledge is itself gained by deference. In other words, the issue is not between agents who have deferred and those who have not, but rather between those whose knowledge that *p* is isolated from knowledge why *p*.²¹

Even if the knowledge/understanding divide is not as wide as it first appeared, the general point of the understanding account seems important. Nevertheless, it doesn't give the whole story. Some negative features of moral deference can remain despite understanding, and some positive features of non-deference can exist without understanding. Consider first:

Sam the Schizoid Ethicist: Sam knows the ins and outs of moral theory. He's read all the stuff, has taken numerous courses in normative and metaethics, and has written not one, but two excellent Ph.D. Dissertations. He defended each of them against the most rigorous examiners who subsequently suggested

publication. The problem is, despite his expertise, he is of two minds about ethics. One of his dissertations was titled “In Defense of Utilitarianism,” and the other was titled, “Against Consequentialism: A Deontological Manifesto.” For just about any moral issue, he can provide a comprehensive explanation and defense of the Utilitarian verdict as well as the Deontological verdict. This is fine, of course, except when the views disagree. Suppose he is confronting an issue on which they disagree—whether, say, to harvest Healthy Harold’s organs to save five others. This is an important decision, so he asks Google Morals “Which is true, utilitarianism or deontology?” It says “utilitarianism” so he judges that he should harvest Harold’s organs.

In Sam we have a case of deference, and assuming utilitarianism is true Sam not only knows that he should harvest Harold’s organs, but he understands why. It’s easy to imagine, in fact, that he can satisfy conditions i-vi as well or better than any of us. But there is still something wrong. Admittedly, not too wrong; as described his situation is somewhat enviable. But another fellow is not so well off:

Urkel, the Ungrounded Ethicist: Urkel, like Sam, has a deep understanding of ethics. But he lacks moral intuitions of any sort. That moment where you see a case and think *hmm, that doesn’t seem right . . .* He doesn’t do that. He also doesn’t have intuitions about the axioms of moral theory, even though he thoroughly understands what a particular set of axioms would commit him to. To solve his problem he consults Google. “Is causing pain unnecessarily a bad thing?” “Yes,” says Google. “Should one hang an innocent man to prevent a race riot?” Google says “No.” Basically, Urkel has the complete moral calculus in place, he just doesn’t have his own inputs. Once he gets them, though, he can justify and explain until the cows come home. He grounds out when it comes to basic judgments, of course, but so do we all.

Urkel has understanding, but he is seriously deficient. His is a case of foundational deference, and foundational deference is, on the face of it, one of the least palatable forms. In some sense, of course, he doesn’t understand his basic judgments.²² But it’s not clear that those of us who rely on moral intuitions are any better off. Yet there seems to be a difference between us, and Urkel seems the lesser agent. Both of these guys, Urkel and Sam, lack some positive features possessed by agents who lack understanding. Consider:

Alastair, the Errant Ethicist: Like Urkel and Sam, Alastair is a well-trained ethicist. He has dedicated no small part of his life to working through problems of ethics, and he has strong intuitions. He is able to provide as good a defense of those intuitions as anyone can, and when it comes to particular cases he has a well worked out calculus that generates judgments in a snap. He can explain his decisions, answer the objections of others, and make people who disagree with him look silly. But he is wrong. Somewhere or other he made, we can suppose,

a subtle philosophical error of the sort to which we are all vulnerable. So, when it comes to a particular judgment, he is in error.

Because of the factivity of understanding, Alastair lacks understanding, while Urkel and Sam have it. There is, of course, no question about who is correct in their judgments, when they differ. By hypothesis, Alastair loses this battle. But there is also a virtue Alastair has that the others don't, and this virtue seems to be precisely what characterizes people who reason their way to moral views in a way that incorporates basic intuitions. This virtue is lacked by people who defer. It is tempting to remove the factivity condition from "understanding" and claim that Alastair lacks this "quasi-understanding." Perhaps there is a way to do this, but it is not obvious. The point of the understanding account is not that the deferring agent lacks an understanding of "p." The deferring agent might understand that proposition fully well. What she lacks is an understanding of "why p" where that can only mean "why p is true." So the factivity of understanding is not incidental. But if this is the case, Alastair lacks understanding. Intuitively, though, he doesn't have the failing that characterizes moral deferrers. This suggests that understanding is not quite what we are looking for. So, I think, does the following case:

Unsophisticated Ursula: Ursula is unlike Alastair, Sam and Urkel. She doesn't have any explicit theories about the morality of her actions, she has no philosophical training, and she's not very good at articulating the reasons for her action. She is also not particularly good at understanding ethical explanations of much sophistication. She does, however, have a very reliable disposition to do the right thing, to feel appropriately in the face of moral situations, and to recommend the right thing to others. When it comes to articulating reasons for her actions, she is likely to defer to those with more training, practice or talent at reason giving.

Ursula lacks understanding. She does not satisfy conditions i, ii, v and vi of the understanding account. As far as the theoretical reasons for her actions go, she defers. As philosophers, we are apt to look down on Ursula, but there is something very admirable about her, something that might well be absent in all three of our other deferrers, and that seems likely to be absent whenever someone blindly defers on moral issues. This again suggests that we must look beyond understanding.²³

Three Non-Epistemic Accounts

Before turning to my own account, it is worth looking at three other tempting options. While they are also inadequate, they reflect intuitions which can be captured in a complete account.

Autonomy

Given the inadequacy of epistemic accounts of the asymmetry of deference, it is tempting to look in the direction of something like autonomy. To the extent that one who defers relies on someone else to know what to do, he might seem to be illicitly dependent on that person. This is one reason, perhaps, why we would be reluctant to place full blame on an agent who deferred but was intentionally misled: they have become a tool for the one who misled them and the blame is allocated accordingly. Even if the source is not being manipulative, the dependence of deferrer upon adviser appears to be in conflict with the deferrer's autonomy.

The topic of autonomy is among the most complicated in philosophy, and perhaps there is a sense of autonomy which will make this analysis work. But it is unclear why deference would prevent an agent from acting autonomously. With respect to which action does the agent lack autonomy? It doesn't seem that the agent is acting unfreely in deferring. Assuming she was an autonomous agent at the time, she could freely and autonomously choose to defer. Similarly, after the agent possesses moral knowledge, it seems she is free to do the right or the wrong thing and there is no reason she is not doing so autonomously. There is, to be sure, a sort of dependence of the deferrer on the adviser, but it doesn't seem that this dependence is in conflict with autonomy. One's actions, intellectual and otherwise, are dependent on many outside factors. If this robs one of autonomy, autonomy is not to be had at all. While it is appropriate in some cases to give the lion's share of praise or blame to the original source of the moral judgment, it's not because the deferrer lacks autonomy.

Acting for the Right Reasons

A related suggestion is that when one defers on moral matters one is not believing or acting for the right reasons. If my students refrain from cheating simply because they believe they will get caught, they are not really acting from the right reason. Similarly, we tend to think that if someone refrains from lying because their priest told them not to, they are acting for the wrong reasons. They simply aren't "in touch" with the reasons why lying is wrong. This suggests that the person who holds a moral belief due to deference is not in a position to act or believe for the right reasons.

While this seems plausible, it makes a basic mistake. One can imagine two individuals learning from their priest that lying is wrong, but Sally might tell the truth because lying is wrong while Bobby might tell the truth because the priest said that lying is wrong. There is a difference between these two, and Sally's position is if anything the more natural. If S believes p based on T's say-so, the subject still believes p—she doesn't simply believe that T reported that p. If the subject comes to know p based on T's say-so, she knows p—she doesn't only know that T reported that p. In fact, S might

not know or believe anything about T at all. So, if deference can provide knowledge of moral facts, then once a subject has properly deferred about the fact that lying is wrong, she knows that lying is wrong. When she then tells the truth in light of that knowledge she does it because she knows lying is wrong. Thus, she's acting for the right reason. Perhaps she knows it is wrong because of testimony, but we shouldn't infer from the fact that it is only by Q that S appreciates the right reasons that Q is (or is a part of) S's reason to act.

Doing One's Job

If I'm to have an operation and my surgeon comes in with a grin and a medical textbook, I'll be concerned. My concern will be mitigated if I know that he is flawless at using the book, but I still won't think him much of a surgeon. It is his job to know that stuff. He might succeed in surgery, not as a surgeon, but rather as a practiced reader with a good pair of hands.

If a theoretical physicist is having doubts about the number of dimensions required by string theory and he asks another physicist, there is something wrong. Even if he comes to know by that route, he is something short of a respectable physicist because he didn't confirm the reasoning. The same is true if, say, a top physicist claims to have solved the problem of quantum gravity and another top physicist simply defers on the matter without checking the work.

The problem with these subjects is that they do not really seem to be doing their job. There is a division of labor, so it's permissible to defer when a matter is outside of one's ambit, but to defer with respect to one's own field is to abandon one's duties. One might get the job done, but one really isn't pulling the weight one has implicitly agreed to pull by taking that job. We expect more of physicians than of physician's assistants, and we expect more of theoretical physicists than we do of physics teachers: they have different jobs and therefore different obligations. The latter can defer about their fields, but the former should not.

It is very tempting to think that the problem with someone who defers on moral matters is really the same as that of these deferring professionals. Just as it is a physicist's job to figure out the truths about (at least her subfield of) physics, it is the moral agent's job to figure out the truth about morality.²⁴ By deferring about moral matters, we might get the job done better than we would have otherwise, but in an important sense we have still shirked the duties that come with that job. This is what seems to be wrong with:

Gary the Googler: Gary wants to do the right thing, he just doesn't trust himself to know what the right thing is. Unlike Urkel, he has intuitions, but he doesn't trust them. He realizes that he could study moral theory, but unlike Sam he hasn't the skill or the patience to work through those convoluted systems of

ethics. And, frankly, Gary fears that even if he did have the patience or the skill, he might wind up like Alastair—sophisticated, convinced and convincing, but still wrong because of a subtle mistake. So, every time Gary comes to a moral issue he defers to Google Morals and then does what it says.²⁵

Gary is troubling, and one might think it is because he has effectively abandoned his job as a moral agent. As appealing as this proposal is, there are concerns. One is that it threatens to push the asymmetry back a step. The proposal really amounts to the suggestion that there is no epistemic division of labor when it comes to moral matters. We can push off the work of knowing and discovering non-normative facts to the relevant experts, but we are each of us responsible for the knowledge and discovery of moral facts. But why should this be? It doesn't help to say that we are all moral agents while only some of us are doctors. All of us have bodies, but we have no hesitation about deferring to others about how to use them or keep them healthy. Furthermore, while it is true that we cannot help but be actors, for even choosing not to act is a type of action, this very fact provides a defense of the moral deferrer that is unavailable to the professional. "Granted," says Gary, "I cannot help but act. Given that I'm the sort of thing that can act, so much is logically true. But I didn't sign up to be a deep moral deliberator and I have doubts about my competence at it. If that is my job, I hereby quit. The least I can do is ensure that I am acting correctly, and for that I choose to defer." To be sure something is wrong with Gary, but it is not the same thing as what is wrong with our professionals. Gary is rather reflective and responsible given his view of his competence, and if professionals adopted a similar perspective and quit their jobs we would have no further gripe. Somehow we don't let Gary off so easily. Thus, it seems, we still have an unexplained asymmetry.

Virtue, Subjective Integration and Moral Character

Neither the epistemic accounts nor the most obvious non-epistemic accounts seem to capture what is at issue in the asymmetry of deference. Most of these accounts get something right, however. The understanding account and the "acting for the right reasons" account draw attention to the fact that the beliefs sustained by deference are oddly isolated from the relevant reasons for the action. The autonomy and "doing one's job" accounts capture the intuition that the deferring agent has a peculiar reliance on others for something that we often want to see stemming from the agent. What we need is an account that captures these intuitions while still issuing the correct verdicts about Sam, Urkel, Ursula and Alastair.

Since it seems implausible that the act of deferring is problematic, or that the resulting actions are lacking, we should remain focused on the agent. I suggest that the problems we see with deference stem from the fact that the

beliefs sustained by deference are largely isolated from the *moral character* of the agent. We admire good agents not simply because they generate good actions but because their actions reflect their virtues. A good deferrer, like Gary, has an important virtue: he is motivated to get the right moral beliefs and is good at doing so. While this is important, the beliefs and actions Gary sustains by deference only reflect *that* virtue, and not, say, the virtue of selflessness, honesty, or whatever virtues give rise to the beliefs, motivations and emotions that are typical for the moral non-deferential agent. The part of Gary's moral character that bears most directly on these issues, if it exists at all, is being bypassed because of his deference. In addition, there is no real guarantee that Gary's Googled belief will properly integrate with the moral character he has.

The relationship between deference, moral character and the virtues is complicated, and to see how it all works we need a basic picture of what is involved in moral virtues. Once that is on the table, we can take a look at the various ways deference can reflect and affect moral character. My proposal doesn't require a particularly idiosyncratic account of virtues, and we needn't get into the details. The view is roughly Aristotelian, and my exposition of it will basically follow that of Julia Annas:

What is it for Jane to be generous? It is not merely that she does a generous action, or has a generous feeling. Either or both could be true without Jane's being generous. . . . For Jane to be Generous, generosity has to be a feature of her—that is, a feature of Jane as a whole, and not just any old feature, but one that is persisting, reliable, and characteristic.²⁶

To have a virtue is to have a reliable disposition to act and feel in certain ways.²⁷ “Further, a virtue is a disposition which is characteristic—that is, the virtuous (or vicious) person is acting in and from character when acting in a kindly, brave or restrained way.”²⁸ A generous person will be disposed to perform generous actions, and will have the feelings appropriate to the situation. They will be reliably motivated to be generous, one suspects they will have deeply seated intuitions about the fact that one should give under the appropriate circumstances, and they will feel generous—in other words, they will feel good about the giving and not, for example, suffer the selfish pangs of loss following every act of “generosity.” The generous person will also no doubt have a host of beliefs, implicit or explicit, about the ways one should act. The sophistication of those beliefs and their relationship differs—and this difference can play an important role in explaining some problems generated by deference—but the beliefs will be present nonetheless.

So, a virtuous person has dispositions to act, feel, and believe in certain ways. But it seems that one can have these dispositions without having the relevant virtues—or at least without having them to the fullest extent—if these virtuous features were not sufficiently integrated. One could, for example,

have a belief about generosity without that belief either stemming from or giving rise to the dispositions and feelings that characterize generosity. The features that characterize a disposition must be unified in an important sense, reflecting a sort of subjective integration on behalf of the virtuous agent.²⁹ This unification between beliefs, feelings and dispositions to act is part of what allows virtues to be dynamic, allowing beliefs to reinforce feelings and dispositions to action, and vice versa.

How does this picture of virtues explain what is sub-optimal about morally deferring agents? I think it can do so in three general ways, roughly corresponding to shortcomings of the agent pre-deference, shortcomings that can be immediate consequences of deference, and difficulties deference introduces for the longer term development of the agent's moral character. More specifically, the virtue account can explain the problems with moral deference in at least the following ways:

- Deference can indicate of a lack of certain virtues, not only because the agent does not know the target fact beforehand, but because the agent lacks either the ability or the drive to get the belief by non-deferential means.
- Deference, and the actions resulting from it, bypass moral character in that the beliefs and actions do not stem from the agent's present virtues (if they are present).
- Deference can result in an agent's having a virtuous belief, and performing a virtuous action, without the agent possessing the relevant virtue.
- Beliefs attained by deference can fail to integrate with the rest of the agent's present beliefs, and can fail to provide the proper ground for new beliefs. This prevents the agent from achieving higher degrees of virtue.
- Deference can either undermine, or frustrate the development of virtues insofar as the agent is likely to be in a poor position to feel or act reliably.

Not every case of deference, of course, will involve all or even most of these shortcomings. But together, these can account for much of our sense of what is wrong with moral deference.

Most clearly, the fact of deference is likely to indicate a deficiency in moral character prior to the deference. This is not, however, the trivial deficiency that stems from fact that the deferring agent didn't know the target fact before the deference. As I noted earlier, this is true of both moral and non-moral deference, and it is true of anyone who gains knowledge in any way whatsoever. Moral beliefs are slightly different in that we often expect it to be the case that they stem from a person's feelings and intuitions which are indicative of a person's moral character. When a person is ignorant about non-moral empirical matters, we are likely to think it merely a matter of ignorance, or perhaps the person's intellectual abilities. If I don't know the directions from New York to Boston, few would find me seriously lacking. Even if a person didn't know whether a circle had angles, we wouldn't judge that person's character. If a person must defer about whether suffering is

wrong, on the other hand, we are inclined to think that this reveals more than a mere gap of knowledge or understanding. Something is wrong with the person. This is, I think, because the person must be lacking intuitions, feelings and other non-cognitive attitudes about suffering, and that a lack of these things indicate the person is lacking virtue. One doesn't reveal these things when one defers about non-moral matters.

The preceding story draws on the fact that sometimes we feel deference is problematic not because of what it creates in the person, but what it indicates. It's perhaps more important that deference can be problematic because of the state it leaves the agent in after deference. For one thing, when one receives moral knowledge by testimony, it is very unlikely that one is receiving the disposition to react emotionally as well. In other, more typical ways of coming to know moral truths, this isn't the case. To see the point, consider a case in aesthetics. Suppose you are told that Verdi's *Aida* is a great musical work. You are rather indifferent to classical music, and to opera in particular, but you recognize that your interlocutor knows her stuff. It is not unusual, or unwise, to defer in a case such as this—you might come to believe that *Aida* is a great musical work, but at the same time you might recognize it doesn't do much for you. Such cases are not particularly worrisome, but there is something odd there. Someone could even become quite good at detecting the earmarks of such great music, but if they lacked any sort of "emotional" or aesthetic reaction to the opera, they fall somewhat short of being a good critic.³⁰ If, however, they come to know that *Aida* is great by, perhaps, studying the history of opera, coming to understand the dramatic and musical developments it reflects, they are significantly more likely to come to an appreciation that might well make them a passable music critic.

The same holds for the moral domain. If I am told that causing suffering is wrong, all things being equal, I might come to know that it is wrong on that basis, but there is no guarantee that I will thereby develop the emotional reaction that most of us have—and most of us think we should have—in the face of suffering. I could still fail to care, in an important sense, that suffering is wrong. Psychopaths, who apparently lack such a reaction, could nonetheless defer to Google Morals and attain justified moral beliefs, though many of them might not.³¹ Similarly, even if I learn that the basic needs of others are as morally significant as my own basic needs, or those of my family and friends, there is no guarantee that I will have any emotional connection to those needs. Even if a psychopath becomes good at preventing needless suffering and does so because he knows that preventing it is a good thing, there is something wrong with him if he remains emotionally indifferent to it. Similarly, if I am plagued by a selfish burning when I give to others, despite the fact that I know it is the right thing to do, I am lacking.

Just as deference on moral matters doesn't give the agent the appropriate emotional responses, it doesn't guarantee that she acquires the appropriate motivations. Just as learning that *Aida* is an excellent opera might not impart

the least desire to listen to it, so one might come to know that it is good to give a great deal of one's wealth to the needy without being inclined to do so.³² Such a disconnect seems incoherent in non-normative matters. (It certainly doesn't make sense for an agent to know that *p* but refrain from believing *p*.) Deference introduces the potential for a disconnect in the normative realm that doesn't exist with respect to the non-normative realm, and such a disconnect is of great significance to moral agency.

Motivations and feelings are likely to be disconnected from beliefs sustained by deference, but so are moral intuitions and related non-cognitive attitudes. Moral judgments are often simpler than the understanding account portrays. Many of us have basic reasons for which we can provide no further explanation. If we are asked why needless suffering is bad, perhaps the more ingenious of us can say something but the rest of us will just insist that it is—that we simply see it is so. What these basic intuitions amount to, where they come from, and whether they provide justification for our beliefs is beside the point. The fact is, we do have these intuitions and they play an important role in our moral reasoning, and even if deference provides moral knowledge it is unlikely to instill the relevant moral intuitions. When Urkel defers about basic matters, there is no guarantee that he will also gain the reserve of intuitions or pro-attitudes that the rest of us have about these matters.

It is tempting to think intuitions only form a foundation for moral reasoning, coming into play when basic moral cases or axioms are under consideration. But this is not the case. Even when we can provide significant explanations for why we hold a particular moral belief, understanding why it is true, it is often the case that the belief is also held in part because of an intuition. A Kantian can be quite articulate about why it is that one cannot harvest someone's organs to save five others, but it would be disingenuous of her to pretend that this explanation does not come alongside a strong intuition that something is wrong in this case. Even if the Kantian explanation is correct and the reasoning behind that explanation is what is sustaining the belief, the intuitions play an important role.

It seems plausible that this is a part of the story lacked by the understanding accounts. Urkel is flying emotionally blind, while Sam accepts an entire system of morality entirely independently of his reactions to cases. Alastair, on the other hand, might be wrong, but his emotive reactions are not alienated from his moral judgments. And Ursula, on this score at least, is likely to be better than all the others despite the fact that she lacks understanding. She has a stable disposition to act and react appropriately, even if she lacks the ability to articulate or even understand the deep reasons behind her actions. She acts out of a set of intuitions, pro-attitudes and feelings that are lacking in many of the others, and the beliefs that she does have about what should be done—simple as they may be—are deeply connected to her dispositions to feel and act. What Ursula has is unlikely to be gained by deference.

Nevertheless, it seems that Ursula is missing something, and the virtue account has the resources to explain what she is missing in a way that captures the intuitions behind the understanding account. We can follow Aristotle in distinguishing between two “levels” of virtue: natural and complete:

For each of us seems to possess his type of character to some extent by nature, since we are just, brave, prone to temperance or have another feature immediately from birth. However, we still search for some other condition as full goodness, and expect to possess these features in another way. For these natural states belong to children and to beasts as well <as to adults>, but without understanding they are evidently harmful. At any rate, this much would seem to be clear: just as a heavy body moving around unable to see suffers a heavy fall because it has no sight, so it is with virtue. <A naturally endowed person without understanding will harm himself> But if someone acquires understanding, he improves in his actions; and the state he now has, though similar <to the natural one> will be virtue to the full extent.³³

Among other things, this passage suggests part of Aristotle’s argument for the unity of virtues. We need not follow him so far, however, to agree that there are different “levels” of virtue. The general picture can be painted in numerous ways, and it doesn’t necessarily depend upon there being virtues imparted by nature.³⁴ The idea is rather that there is a level of virtue that is associated with the dispositions to feel and act appropriately, as well as to have certain intuitions and pro-attitudes. Someone like Forrest Gump might have this level of virtue. Gumpish virtue is not optimal, however. One wants not only to act in a certain way but to understand why one is acting in that way in the relevant circumstances. This allows one to determine how to act in new, slightly different cases, it helps one to shape one’s dispositions in response to changing circumstances, and it enables one to be a better teacher of morals to others, such as one’s children.

It is plausible that Ursula is lacking this more complete virtue, and it is also likely—as the understanding account argues—that it is lacked by one who defers on moral matters. If someone doesn’t know why, say, first term abortion is morally permissible, they will not necessarily know whether infanticide is morally permissible. They will be frustrated by new cases, and cannot be relied upon to apply a general understanding of the situation. They will also be less reliable as teachers of morals to others. Thus, moral deference tends not to inculcate or reflect a natural virtue, but it also frustrates the attempt to achieve a complete, more reflective virtue. Thus, even if the dispositions somehow came with the deference, the agent would remain in a sub-optimal position.

What about Alastair? Something seemed good about him—better than Ursula, for example—even though he lacked understanding (because his beliefs were false). Here I am inclined to deviate a bit from the Aristotelian

line. While a deferential agent is unable to provide the actual reasons why something is good or bad, this is something many of us—even the most reflective of us—cannot do. It is particularly worrisome, nevertheless, if I am also unable to explain why my belief appears to be true in light of my own reasons. This is what Alastair is able to do, and this is part of what we expect of people when they disagree with us. The fact that we expect people to do this when they disagree with us, and thus when we think they are wrong, suggests that our expectations do not depend upon our interlocutors providing the right reasons. If someone is unable to provide a justification of their moral claims in terms of what they take to be relevant moral reasons—rather than by pointing to an authority—we feel they have failed us. The failure is no doubt greater if they are actually wrong, but it is a failure either way. Moral deference all but insures such failure.

This failure is a result of the fact that the belief attained by deference is cognitively isolated and is not introduced in a way that guarantees integration and coherence for the subject. This lack of cognitive integration poses a further problem given the way people tend to engage in moral reasoning. Few of us reason up from a set of basic moral axioms. Instead, moral reasoning tends to be more of a holistic affair where we not only embrace the consequences of our previously accepted moral principles, but go back and revise our moral principles in light of unforeseen repugnant consequences. We strive for a sort of reflective equilibrium wherein our judgments “accord with our considered convictions, at all levels of generality, on due reflection.”³⁵ To the degree that deference introduces cognitive isolation, it will make it more difficult to achieve this reflective equilibrium. This has several drawbacks. First, it seems to have an epistemic drawback since the deferring agent will be less able to achieve coherence among his beliefs. While he will be able to detect obvious contradictions among his beliefs, he will not be in a position to appreciate more subtle forms of incoherence. He will be less sensitive to the fact, for example, that his new belief is only supported by reasons which would provide defeaters for some of his other beliefs. And because he might not recognize that his belief commits him to other beliefs, because he might not understand how to apply his belief to new cases, he might not see how those other beliefs would be in tension with beliefs he already has. Second, making moral judgments will presumably involve acting, as well as making demands on oneself and others. Someone who is a disunified moral agent will be more likely to act and make demands that do not reflect a coherent moral view and that therefore appear to be random and spurious. This makes them unpredictable both in how they behave and in how they dispense praise and blame. These facts make it difficult to trust such agents and it consequently make civic and moral relations with them problematic.³⁶

Just as portions of the virtue account can capture what we get from the understanding account, part of the virtue account can explain some of the intuitions behind the “doing one’s job” account. That account was driven by

the idea that there are some non-moral domains where deference seems just as weird as in the moral case. If a doctor deferred (too much) about a key medical procedure, or if a physicist deferred about string theory, we would think something was strange. We would think, it was suggested, that they weren't doing their job. In fact we now have a slightly different diagnosis, which acknowledges this intuition while still maintaining there is a difference between such cases and the moral case. Virtues, it is often pointed out, are much like skills.³⁷ One is not a skillful tennis player by executing a tricky shot, but by hitting a tricky shot from a disposition to hit such shots reliably under the correct circumstances. Furthermore, a fully skilled tennis player understands why that is the shot to hit, and why a similar shot is not the one to hit, and she is disposed to hit the right shot at the right time. Skills, like virtues, cannot be acquired by deference alone, and if one acts out of deference one is not acting in such a way that it reflects one's skill. The doctor who must constantly submit a list of symptoms to a computer is not employing the skill that makes someone a good doctor. Like Gary, he is employing a skill—he is good at looking things up—but it is not the one we ask doctors to have. So the problem is not really that he is not doing his job, it is that he is not doing it in a way that reflects his skill, and we choose to go to doctors instead of librarians for our health problems because they have that particular skill.

Though the analogy with skills is informative, and helps us see the truth in the “do one's job” account, it should not be taken so far as to remove the sense of asymmetry. There are important disanalogies between virtues and skills. For one thing, nothing about having a skill requires that one have a particular feeling when exercising it. As we have seen, having a virtue does require this—one is not selfless if one gives money to others but feels terribly selfish pangs at doing so. Since such feelings are not likely to come simply because one has a deferential moral belief, deference is particularly problematic in the moral domain. Relatedly, skills do not really involve a commitment to the value of performing the skillful action. Serena Williams, apparently, would rather shop than play tennis, but she is no less skillful because of it.³⁸ If someone would rather shop than save a drowning child, no amount of skill would make her virtuous. Since deference with respect to moral beliefs doesn't necessarily provide one with motivations or pro-attitudes, this is a further reason deference is particularly problematic in the moral domain.

We have seen several virtue-related reasons why the agent who defers on moral matters will be sup-optimal. The belief of the agent and the actions it gives rise to will not necessarily redound to the agent's character because they don't stem from a virtue of the agent. Because the deference fails to bring with it feelings, intuitions and motivations, the resulting belief does not give rise to a virtue or reinforce existing virtues in the way the belief might if it were gained in other ways. Because these features are lacking,

the agent is not guaranteed to have the relevant motivations, so there is little guarantee that the agent will be a reliable ethical actor, and because the agent lacks understanding of why the belief is true he will be less able to apply the belief to new, slightly different cases. Finally, because integration of the belief into his own reasoning about the moral domain is not guaranteed, the belief will be too isolated to affect the agent's other virtues and beliefs, and the agent will be unable to provide his reasons for his beliefs, or develop a more complete, integrated type of virtue.

There are apt to be numerous objections to this virtue-based account. One is simply that it might seem to presuppose a contentious "virtue ethics." In fact, this account of the asymmetry of deference does not presuppose "virtue ethics" in any serious way. All it presupposes is that there are such things as virtue and moral character and that we care about them. We can remain agnostic on many issues such as whether or not virtues are analytically more fundamental than good actions, or whether good actions are good only in a derivative sense.

Nevertheless, one might object that this account makes feelings and emotions too important to moral agency. We can easily imagine cases where individuals act morally but are relatively affectless, perhaps because they are depressed. Kant certainly seems to have thought so, and might even have thought the lack of affections would make for a moral advantage. He famously wrote:

But assume that the mind of that friend to mankind was clouded by sorrow of his own which extinguished all sympathy with others, and though he still had the power to benefit others in distress their need left him untouched because he was preoccupied with his own. Now suppose him to tear himself, unsolicited by inclination, out of his dead insensibility and to do this action only from duty and without any inclination—then, for the first time his action has genuine moral worth.³⁹

Must the subjective integration account disagree?

Clearly the virtue account is more Aristotelian than Kantian, and I wouldn't be too dismayed if I had to side with the Greek. In fact, though, this account needn't disagree with much of the Kantian intuition. It is important to recall that the virtue account does not claim that the deferring agent's *actions* are immoral or even sub-optimal. The claim is that the *agent* is sub-optimal. It is thus open for me to agree with Kant that the action is moral because it was performed out of duty and not out of emotion. I only disagree with Kant if he maintains that there is nothing sub-optimal about the affectless agent. While Kant might seem to suggest as much in the Groundwork, I see little pull to that extreme view. At the very least, emotions and feelings tend to reinforce our motivations to act, and though it might be ideal for the affectless agent to drive himself to action by the inspirations of

duty, it is almost certainly the case that the agent with affections will more reliably act in accordance with them. Kant himself recognizes the importance of feelings such as this in the *Metaphysics of Morals*.

But while it is not in itself a duty to share the sufferings (as well as the joys) of others, it is a duty to sympathize actively in their fate; and to this end it is therefore an indirect duty to cultivate the compassionate natural (aesthetic) feelings in us, and to make use of them as so many means to sympathy based on moral principles and the feeling appropriate to them For this [compassion] is still one of the impulses that nature has implanted in us to do what the representation of duty alone might not accomplish.⁴⁰

Although I myself would be inclined to give more role to emotions and feeling than Kant, one needn't reject the Kantian intuitions to accept the relevance of emotions and motivations to moral agency.

One might grant all of this but think that emotional and motivational integration cannot be what is at the heart of the asymmetry of deference because there are plenty of ethical beliefs which do not tie into emotions and motives, and there are plenty of things that do not tie into morality that are tied into emotions and motives. Suppose, for example, I deferred to Google *Morals* on the claim that "Eskimo infanticide is less wrong than infanticide in other cultures."⁴¹ Since I am not an Eskimo, there is no need for me to have particular motivations. Furthermore, it's not clear that I need to have emotional reactions here since I will never encounter the relevant cases. But there's still something sub-optimal about my deference, right? How does the emotions and motivations story account for this?

First, even if this objection was granted, the virtue account could still accommodate the asymmetry of deference because motivations and emotions are not all there is to possessing virtues. Complete virtue requires integration of reasons, and someone who believes that Eskimo infanticide is more permissible than other infanticide will lack integration to the extent that they do not know what it is that makes these cases permissible. They do not know, for example, whether this is because of particular exigencies faced by the Eskimo or whether it is because Eskimos are less than human. If an opportunity to interact with Eskimos did arise, this ignorance could be quite significant. Further, the lack of emotional and motivational integration is problematic even if it didn't. For one thing, despite the new belief, one's feeling of outrage at the Eskimo parents might not dissipate, and this may well prevent one from reacting appropriately to anti-Eskimo racism. For another, it is certainly odd that someone could have an emotional reaction to infanticide in general but then suddenly defer about Eskimo infanticide. Even if the deference is warranted, surely this individual is somewhat internally fragmented if they do not know what distinguishes the different cases of infanticide.

But aren't there non-moral cases that are tied to emotions and motivations in the same way? It is odd to defer about someone's sexual attractiveness, for example, and it might seem it is because the resulting judgment is disconnected from one's emotions and subjective responses. But if so, one might think my explanation does not really capture the appropriate asymmetry.⁴²

In fact, I am happy to admit that a portion of my account applies to aesthetic domain, and at least some judgments of attractiveness are of this sort.⁴³ This is still a normative domain and, as I maintained earlier, there seems to be something fishy about deferring in any such domain. Nevertheless, I must agree that there is something different about such cases, but these differences wind up fitting nicely with the virtue account. For one thing, we simply tend to care more about the moral virtues. Aside from that, in the moral domain we are inclined to think there are principles to which our emotions have to answer and there is some need to square our reactions with those principles. It is not clear there are similar principles in the aesthetic domain. This would mean that the lack of integration in the ethical domain should be more worrisome, since one's other emotional commitments implicitly commit one to other emotional reactions in the case of ethics, but not in the case of aesthetics.⁴⁴

Finally, one might object that my account seems to presuppose that virtue is inconsistent with deference, but this presupposition is patently false. After all, some of us have virtues we were not born with, and no doubt that has a lot to do with our deferring a great deal as children. Deference as a child seems to be a precondition for attaining virtue, in fact. On the face of it, this is inconsistent with my claim that deference frustrates the attainment of virtue and undermines the ascription of moral beliefs and actions to moral character.

Nothing I have said is inconsistent with the claim that some moral deference is necessary if one is to become a moral agent. In fact, nothing about this account says that one shouldn't defer often. This account only maintains that one who has a belief by deference is less optimal (all things being equal) than one who gains the belief in a way that doesn't involve deference. It is unquestionable that children often do, and should, defer. We think nothing is particularly problematic about a child who defers because we cannot expect a child to get at certain truths in any other way. Expecting them to reason to moral conclusions all the time would be absurd. This does not mean, however, that the child is not sub-optimal in part because of the reliance on deference.

In fact, I suspect deference really plays much less of a role in our moral development than it might seem. Some of a child's beliefs are sustained just on the say-so of others, but the behavior of children is reinforced and undermined in many other ways. We do not raise children by testifying to them about the norms of the world. We shape them by encouraging some behaviors and discouraging others, in hopes that they acquire the right sorts

of drives and feelings to allow them to be virtuous agents. We would be disappointed if their moral beliefs continued to be sustained by deference and thereby failed to be integrated into their characters.

When we think of someone gaining virtues by deference we must be clear about what we are imagining. Deference requires not just that the agent come by a belief by discovering the belief of another agent. It requires that the deferrer's belief is based on the other's belief and sustained because of the fact that the other agent held it. It is in fact difficult, perhaps impossible, to imagine someone who has a belief that is subjectively integrated to the point that it stems from a virtue, but still has it because of deference. This would require, after all, that the belief be integrated with her intuitions, motivations and other values. It would require that she knows not only that *p*, but why *p*. It is doubtful that in such a case her belief in *p* would really be based on the say-so of someone else given its rich support within her own character. If the belief were sustained by the deference, it would not be cognitively integrated, which is required for complete virtue. Integration is at least in part a causal matter, involving relations of basing and dependence. For a belief to be integrated into a subject's cognitive and conative economy, it must be sensitive to changes elsewhere in the economy. This would not be the case if the belief were truly deferential.

Conclusion

As I noted at the beginning, there are many different ways in which one can defer in moral matters. There is empirical deference, derivative deference both with and without ignorance, and foundational deference. I've now offered one suggestion about why moral deference is problematic—deferring about moral matters results in judgments which are not subjectively integrated, and so neither stem from a subject's virtues, nor are likely to reflect or further the subject's moral character, so long as the belief is sustained by deference. In addition, the deferring agent is likely to be inconsistent in applying her knowledge because she will have more difficulty achieving coherence, and she will be unable to account for her actions.

Of course integration comes in degrees, so it's not really the case that someone either has moral character or doesn't. We recognize this when we say of an adolescent that his character is not yet fully formed. It's also not the case that after deference the resulting actions simply bypass character. Whether or not an action is creditable to an agent's character is itself a matter of degree. One advantage of this diagnosis and the fact that it admits of degree is that it fits nicely with an intuitive ordering of the perniciousness of the different types of deference. In particular, derivative forms of deference are probably less likely to result in complete incoherence since the belief based on deference is ultimately suited to the subject's value set, as well as the appropriate feelings and intuitions. Since the judgment is not based on

reasoning from that value set, the subject will still be poor at articulating her reasons in detail and at seeing connections and contradictions in her reasons, and so she will have trouble achieving complete virtue, but at least her more basic emotions and reasons will be in harmony with the judgment based on deference.

This analysis also has the benefit of accounting for what is odd about self-deference. From the fact that a belief or value was (or will be) be integrated with one's virtues, it doesn't follow that it would be at the moment of deference. If one was a moral amnesiac, and deferred to one's past belief that abortion was wrong, that belief would be subjectively isolated. If one discovered that in the future one would believe that abortion was wrong, one would have reason to expect that the contrary belief would eventually lack integration, but there is no reason to expect that the pro-life belief would enjoy integration now. Our moral characters evolve and it is inappropriate to try to force a value upon them at one time just because it might be a value at another.

The connections between virtues and moral deference provide us with an impressive fund of resources to explain why it is we feel there is something wrong with using a Google Morals, while there is nothing particularly problematic about using Google Maps. What is perhaps most appealing, however, is that this account explains the asymmetry in terms of something squarely within the moral domain. There might be epistemic dangers associated with moral deference, but the real harm is the crippling effect such deference can have on the moral character of the deferring agents.⁴⁵

Notes

¹ I owe both the term "asymmetry of deference" and my first exposure to this topic to David Estlund, who first circulates an unpublished paper on the topic in the late 90s.

² I don't mean to imply that non-cognitivism, subjectivism, anti-realism or moral skepticism provide explanations of the apparent asymmetry of deference. At the very least, these views do not automatically provide explanations of our intuitions, as McGrath (2009) points out.

³ Compare: "It would have been preferable to use a phillip's head screwdriver, but since one was not available you were right to use the flathead."

⁴ This is not to say that there aren't some cases where the act can be faulted for resulting in a problem in the agent. It's just to say that the ultimate problem is in the agent.

⁵ Thanks to Joshua Schechter for helping me see this important distinction.

⁶ This differs from Driver (2006) for example.

⁷ The Pure/Impure distinction comes from McGrath (2009).

⁸ Some of these distinctions came out of conversations with Walter Sinnott-Armstrong.

⁹ There might be further important differences in this camp, such as whether the adviser knows he is using your values, or whether it is purely accidental that the advice is based on values you hold.

¹⁰ Though there are differences in their views, this sort of account can be found in Jones (1999), Driver (2006), and McGrath (2009).

¹¹ This account is close to that offered in Driver (2006). One finds a little more worry about manipulation in Jones (1999).

¹² Although the movie *Memento* provides an interesting case of self-manipulation in precisely this way. That is surely an unusual case, however.

¹³ For views according to which stakes raise standards for knowledge, see Fantl and McGrath (2002) and Stanley (2005).

¹⁴ Driver (2006) hints at this type of argument and McGrath (2009) develops it further.

¹⁵ Those who do think it is basic include Reid (1764), Coady (1992) and Goldberg (2006) and those who don't include Hume (1739), Fricker (1987) and Van Cleve (2006). For an excellent survey of this literature, see Lackey (2006) and the essays in Lackey and Sosa (2006). Even if one doesn't think testimony is a basic source of knowledge, one might deny that the reliability of a particular testimonial source must be ascertained if it is to provide knowledge. We might have other general evidence of the reliability of testimony that underwrites our acceptance of it.

¹⁶ It is not entirely clear how positions on testimony map onto this issue since deference does not require testimony (as the fMRI case shows). But one can expect that most non-reductionists about testimony would not require that we ascertain the reliability of our sources. Even non-skeptical reductionists should have no problem with what I am arguing here. All that matters for my purposes is that we can take testimony without having to establish the reliability of our particular source.

¹⁷ One has to be careful about just how lucky this subject is to defer to an expert. We can suppose that this subject is in a situation where it is unlikely that he would defer to a non-expert (or someone who was on the wrong side of the controversial matter). So he is lucky to be in that situation, but given his situation he is not lucky to have gotten the truth by deference. For discussions of the different roles of epistemic luck, see Sosa (2007) and Pritchard (2005).

¹⁸ Mark Heller has suggested that our doubts about this agent stem from the fact that we, those judging the subject's situation, know there is a disagreement and can't tell who the experts are, and this affects our judgment of the individual who happens to be in the felicitous situation. Even if our judgment isn't true (as it might turn out on certain forms of contextualism) our negative opinion would be explained by our own awareness of disagreement. It seems, however, that even if this explains judgments in a non-theoretical context, it doesn't explain our judgment that something remains problematic even stipulating that the subject is deferring to an expert.

¹⁹ Accounts of this sort can be found in Hills (2009), Hopkins (2007) and Nickel (2001).

²⁰ Thanks to Jennifer Lackey for helping me realize this point.

²¹ This is similar to Sosa's point (roughly put) that the epistemic externalist can accommodate the reflective knowledge desired by the internalist. The fact that this reflective knowledge is itself grounded in something externalist is not a mark against it. See the essays in Sosa (1991), as well as Sosa (2007).

²² There seem to be two different sorts of cases of deference for someone like Urkel. He could defer about particular cases when he sees or hears about them, or he could defer about moral axioms. The former are foundational in the sense that he might make the judgments without appealing to any other beliefs, while the latter are foundational in the sense that they are about the foundations of morality.

²³ This sort of point is emphasized by Arpaly (2003). I must thank Nomi Arpaly and Peter Railton for leading me to think more deeply about such cases.

²⁴ While I have been attracted to this view for some time, in part inspired by unpublished work by David Estlund, thanks to Joshua Schechter for helping me recognize more of its appeal.

²⁵ There is something incoherent, perhaps, in supposing that Gary defers in every moral matter. For one thing, it seems the decision to defer is a moral matter. For another, it would seem that Gary would need to be able to know a moral matter when he saw one, which requires a sophisticated moral sensibility. Nevertheless, we can imagine a version of Gary who defers as much as possible and this will serve our purpose.

²⁶ Annas (2011) p.8.

²⁷ Annas (2011) p.8–9.

²⁸ Annas (2011) p.9.

²⁹ Annas (2011) p.8.

³⁰ A referee for this journal asks “Suppose that John believes *Aida* is good, but doesn’t enjoy listening to it. However, he does feel regretful about about not enjoying listening to it. Would that satisfy your conditions?” The answer is no, his having these regretful emotions does not make him any better as a music critic, but his having these emotions does satisfy at least one of the conditions for being an agent who values being a good critic. Quite plausibly this gives him a virtue, since valuing a virtue is virtuous, but it isn’t the virtue we are looking for. (See Hurka 2003).

³¹ It doesn’t seem particularly relevant that psychopaths might also, as a matter of fact, not defer. The point is simply that if they did, they would be lacking something. For a discussion of psychopaths, and others who apparently lack the relevant emotional connections despite their understanding of the rules, see Gazzaniga (2008) Chapter 4, as well as Damasio (1994).

³² This claim seems to presuppose motivational externalism, but at least some forms of internalism would allow a similar point to be made.

³³ Aristotle (2005), p.874 114b4–13

³⁴ It has been used to good effect in epistemology by Ernest Sosa, for example. My account owes much to his works in Sosa (2007), (2009) and elsewhere.

³⁵ Rawls (1996) p.8.

³⁶ These remarks draw some inspiration from Dworkin’s remarks on Integrity in Dworkin (1986). Thanks to Luke Robinson for drawing my attention to this parallel.

³⁷ Aristotle (2005), Sosa (2007), Annas (2011).

³⁸ <http://www.dailymail.co.uk/news/article-2081467/Serena-Williams-says-prefers-shopping-sports.html?ito=feeds-newsxml>.

³⁹ Kant (1785) p.252.

⁴⁰ Kant (1797) (MM 6:457).

⁴¹ This example is from an anonymous referee of this journal.

⁴² Thanks to a referee of this journal for leading me to think about this case.

⁴³ There is something odd about deferring that someone is sexually attractive, and this is because of isolation from one’s own tastes and inclinations. But this case is quite different since the isolation is from the reactions that make them true. It would be quite odd if someone found someone worth deferring to in this case, given that there is surely some privileged access we have to our own reactions. In the ethical case the problem with deference seems to remain even if ethical judgments are objective, and the emotional and motivational isolation here relates to the possibility of not responding appropriately to those objective truths. So, there remains a difference between the moral case and the case of sexual attractiveness.

⁴⁴ Thanks to Steve Sverdluk for this suggestion.

⁴⁵ This work owes a great deal to the participants in the Brown University/Princeton University workshop on moral expertise in the Spring of 2011, as well as to the participants in the ROME conference in the summer of 2011. In particular, I wish to thank David Estlund, Elizabeth Harman, Mark Heller, Allison Hills, Clayton Littlejohn, Errol Lord, Sarah McGrath, Joshua Schechter, Paulina Sliwa, and Brian Talbot. The SMU philosophy department was also a great help, especially Doug Ehring, Justin Fisher, Luke Robinson and Steve Sverdluk.

References

- Annas, Julia. (2011). *Intelligent Virtue*. Oxford: Oxford University Press.
- Aristotle (2005). *Nicomachean Ethics*. In Cohen, Curd and Reeve. (eds). *Readings in Ancient Greek Philosophy*. 3rd Edition. Indianapolis: Hackett.
- Coady, C.A.J. (1992). *Testimony: A Philosophical Study*. Oxford: Clarendon Press.
- DaMasio, Antonio. (2010). *Self Comes to Mind*. New York: Pantheon.
- Driver, Julia. (2006). “Autonomy and the Asymmetry Problem for Moral Expertise,” *Philosophical Studies* 128:619–644.

- Dworkin, Ronald. (1986). *Law's Empire*. Cambridge: The Belknap Press of Harvard University Press.
- Estlund, David. (Unpublished) "The Asymmetry of Deference"
- Fantl, Jeremy. and Matthew McGrath (2002). "Evidence, Pragmatics and Justification," *The Philosophical Review* 111: 67–94
- Fricke, Elizabeth. (1987). "The Epistemology of Testimony," *Proceedings of the Aristotelian Society Supplement* 61:57–83.
- Gazzaniga, Michael. (2009). *Human*. New York: Harper Collins.
- Goldberg, Sanford. (2006). "Reductionism and the Distinctiveness of Testimonial Knowledge," in Lackey and Sosa 2006.
- Hills, Alison. (2009). "Moral Testimony and Moral Epistemology." *Ethics*. Vol. 120, No.1, pp. 94–127.
- Hopkins, Robert. (2007). "What is Wrong with Moral Testimony," *Philosophy and Phenomenological Research* 74: 611–634.
- Hurka, Thomas. (2003). *Virtue, Vice and Value*. Oxford: Oxford University Press.
- Jones, Karen. (1999). "Second-Hand Moral Knowledge." *The Journal of Philosophy*, Vol. 96, No.2, pp. 55–78.
- Kant, Immanuel. (1785). *Groundwork of the Metaphysics of Morals*. In *Kant: Selections*. Ed and trans. Beck, Lewis White. New York: Scribner.
- (1797) *The Metaphysics of Morals*. In Gregor, Mary. (trans., ed.) (1996). *Practical Philosophy*. Cambridge: Cambridge University Press.
- Lackey, Jennifer. (2006). "It Takes Two to Tango: Beyond Reductionism and Non-Reductionism in the Epistemology of Testimony," in Lackey and Sosa 2006.
- Lackey, Jennifer., and Ernest Sosa, eds. (2006). *The Epistemology of Testimony*. Oxford: Oxford University Press.
- Nickel, Philip. (2001). "Moral Testimony and Its Authority." *Ethical Theory and Moral Practice* 4: 253–266.
- McGrath, Sarah. (2009). "The Puzzle of Pure Moral Deference," *Philosophical Perspectives* 23: 321–344.
- Pritchard, Duncan. (2005). *Epistemic Luck*. Oxford: Oxford University Press.
- Rawls, John. (1996). *Political Liberalism*. New York: Columbia University Press.
- Reid, Thomas. (1764). *An Inquiry into the Human Mind on the Principles of Common Sense*.
- Sosa, Ernest. (1991). *Knowledge in Perspective*. Cambridge: Cambridge University Press.
- (2007). *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, Vol. I. Oxford: Oxford University Press.
- (2009). *Reflective Knowledge: Apt Belief and Reflective Knowledge*, Vol. II Oxford: Oxford University Press.
- Stanley, Jason. (2005). *Knowledge and Practical Interests*. Oxford: Oxford University Press.
- Strawson, P.F. (1962). "Freedom and Resentment" *Proceedings of the British Academy*, 48: 1–25
- Van Cleve, James. (2006). "Reid on the Credit of Human Testimony," in Lackey and Sosa 2006.
- Wallace, R. Jay. (1996). *Responsibility and the Moral Sentiments*. Cambridge: Harvard University Press.