

The Modesty of the Moral Point of View¹

Karl Schafer

3/5/15

In *Weighing Reasons*, E. Lord & B. Maguire (eds.) OUP

In recent years, several philosophers - including Joshua Gert, Douglas Portmore, and Elizabeth Harman - have argued that there is a sense in which morality itself does not treat moral reasons as consistently overriding.² My aim in the present essay is to develop and extend this idea from a somewhat different perspective. In doing so, I offer an alternative way of formalizing the idea that morality is modest about the weight of moral reasons in this way, thereby making more explicit the connections between this thesis and similar issues in the epistemic sphere. In addition, I discuss how these ideas can transform our thinking about familiar questions in ethics such as the nature of self-effacement, the significance of reflective endorsement, the weight that moral reasons ought to be given in all things consideration, and the plausibility of “indirect” moral theories. Finally, I show that these ideas are compatible even with pictures of morality – such as Kant’s – on which morality might seem to anything but modest about its own importance. In doing so, I stress that it is possible to see morality as modest about the weight of specifically moral reasons, while also seeing all practical reasons as grounded in morality more indirectly – namely, by seeing morality as determining the weight that both moral and non-moral considerations deserve to have in all things considered deliberation.

Modest Evaluative Methods

Suppose you are looking for an evaluative method to trust.³ By an evaluative method, I mean a systematic way of determining one’s preferences over the different possibilities open to one.⁴ At

¹ The ideas in this paper benefited greatly from repeated rounds of feedback from Barry Maguire. In developing these ideas, I was also helped by comments from Errol Lord, Kieran Setiya, Robert Steel, and Jack Woods. The work involved in writing this paper was partially supported by the Alexander Humboldt Foundation.

² See Gert’s “Normative Strength and the Balance of Reasons.” *Philosophical Review* 1:16 (2007): 533-562; Portmore’s “Are Moral Reasons Morally Overriding?” *Ethical Theory and Moral Practice*, 11:4 (Aug., 2008), 369-388; and Harman’s “Understanding Some Cases of Altruism as Permissible Moral Mistakes” (work in progress). All of these essays are more concerned than I am here with *arguing* for the thesis that morality is modest about the weight of moral reasons. Thus, although I will discuss some reasons for finding this idea attractive, my primary concern will be with making better sense of it and its implications. I hope the discussion to follow will complement their important work.

³ Readers familiar with Lewis’ classic paper “Immodest Inductive Methods” (*Philosophy of Science*, 38 (1971), 54–63) should find the next two paragraphs more than vaguely familiar. There Lewis defines the notion of an immodest inductive method as follows:

Let us say that an inductive method C recommends an inductive method C' if the C-mean estimate of the accuracy of C' is not exceeded by the C-mean estimate of the accuracy of any rival method. An inductive method might or might not recommend itself. If it does, let us call it immodest. When asked which method has the best estimated accuracy, the immodest method answers: "I have."

⁴ Thus, an “evaluative method” in this sense is meant to capture normative, deontic, *and* evaluative information.

least for the moment, we can represent such a method by a function C from propositions to real numbers in the interval $[-\infty, \infty]$. You trust a method if your preference function assigns the value $C(p)$ to every p .⁵

One thing you can do with an evaluative method C is to evaluate which evaluative method it would be preferable for you to trust in your current situation. Let us say that an evaluative method C recommends that I trust an evaluative method C' if the value that C assigns to the proposition that I trust C' is not exceeded by the value that C assigns to the proposition that I trust any rival method. An evaluative method might or might not recommend itself to me in my current circumstances. If it does, let us call it immodest in those circumstances. When asked which method it would be best for me to trust, the immodest method answers: "I am."

On the other hand, an evaluative method that is modest is one that answers this question in the negative. That is, it is a method C of which the following is true of me (given my current situation):

There is an evaluative method $D \neq C$, such that $C(\text{I trust } C) < C(\text{I trust } D)$.⁶

Plainly, some evaluative methods will be modest in some situations and immodest in others. But the methods that will be of interest to us here are ones that are systematically modest under normal conditions. Thus, they are methods C of which the following is true:

For any individual (at a time) X under normal conditions, there is an evaluative method $D \neq C$, such that $C(X \text{ trusts } C) < C(X \text{ trusts } D)$.

To capture this feature of these methods, we can define an *über-method* to be a function from centered possible worlds (or individuals at a time) to evaluative methods.⁷ Then the following will be true of any method C that is systematically modest in the manner just described:

There is an über-method D^* , such that for any individual (at a time) X under normal conditions, $D^*(X) \neq C$ and $C(X \text{ trusts } C) < C(X \text{ trusts } D^*(X))$.

Of course, for obvious reasons, the modesty of a method will be most significant when it applies to me under my actual circumstances. Thus, I will sometimes focus on this case for simplicity of exposition. But the modest methods I will be focusing on will all be *systematically* modest in the way just described.⁸

Very often when the distinction between modest and immodest methods is discussed, it is assumed that modesty should be avoided, if at all possible. For example, in Lewis' classic discussion of

⁵ So it takes rather little to "trust" a method in this sense.

⁶ Strictly speaking, to get real modesty, we should also require that there be a *genuine* conflict between the verdicts of C and D in at least some cases. Thus, D should not merely be a way of making C more precise. In what follows, I will generally leave this complication to the side.

⁷ Compare Maria Lasonen-Aarnio. "Single Premise Deduction and Risk." *Philosophical Studies*, 141 (2008), 157–173

⁸ Or at least close enough to being systematically immodest that the difference is of no significance. Nonetheless, it would perhaps be best to say that these methods are modest *almost everywhere* within the space of normal possibilities.

immodest inductive methods, it is taken for granted that immodest inductive methods are preferable to modest ones. And in contemporary discussions of “reflective endorsement” within metaethics, it is often assumed that an evaluative method, faculty, or standpoint can only have normative authority if it is immodest in its own self-evaluation.⁹

There is something natural about these claims. After all, a modest epistemic method is one that calls into question its own reliability. And a modest evaluative method is one that calls into question whether it should guide our choices and actions. Such results are often taken to mean that these methods are self-undermining in a deeply problematic sense. So, for example, it is sometimes suggested that immodest epistemic methods are somehow inconsistent. And, as we will discuss in a moment, the modesty of an evaluative method is often taken to mean that it (or the theory it formalizes) is objectionably “self-effacing”.

This is one way of interpreting the fact that a method is modest, but it is not the only one. For this formal feature of a method might also mean that it (or the theory it formalizes) regards itself as only *part* of the full story about what a rational agent ought to believe or do. For example, the possibility that an evaluative method is modest may simply indicate that it (or the theory it formalizes) recognizes that it does not capture *all* of the considerations to which all things considered deliberation ought to be responsive. In this case, the method’s recommendation of some method other than itself may simply represent its (or the theory it formalizes’) recognition that all things considered deliberation should be responsive to reasons other than those the method formalizes.

Such a method would be modest about its own significance in a much more positive sense of the word. Indeed, depending on the role this method is meant to fulfill, such modesty need not be any embarrassment to it. Of course, it is true that compliance with a modest epistemic method can only be rational when this method is qualified or restricted so as to remove the ways in which it appears to be unreliable, even by its own lights.¹⁰ And, as we will discuss further, compliance with a modest evaluative method will often only be rational insofar as its concerns are held in check by the other sources of reasons for action that this method *itself* recognizes in some sense. So modest epistemic and evaluative methods do seem to be poorly suited to play the role of providing us with a *complete* account of the nature of rational belief formation or choice. Thus, if we are searching for a method that will play *this* role, the fact that the method we are considering is modest should, I think, come as a disappointment.

But this is not the only reason why we might be interested in an evaluative method. For example, we might be interested in this method as a way of formalizing an important subclass of practical reasons or values. In this case, we would not expect this evaluative method to represent the complete story about matters of practical significance. And given this, we should hardly be concerned to discover that this method is modest. Indeed, we might well be pleased with this

⁹ See, in particular, Christine Korsgaard’s *The Sources of Normativity*. (Cambridge: Cambridge University Press, 1996).

¹⁰ In addition to Lewis’ “Immodest Inductive Methods”, see Harty Field. “A Priority as an Evaluative Notion.” in P. Boghossian and C. Peacocke (eds.), *New essays on the a priori*. (Oxford: Oxford University Press, 2000); and Adam Elga. “How to Disagree about How to Disagree.” in R. Feldman and T. Warfield (eds.), *Disagreement*. (Oxford: Oxford University Press, 2010). But contrast David Christensen. “Does Murphy’s Law Apply in Epistemology? Self-doubt and Rational Ideals.” *Oxford Studies in Epistemology*, 2 (2008), 3–31; and Joshua Schechter. “Rational Self-Doubt and the Failure of Closure.” forthcoming in *Philosophical Studies*.

discovery. For in this case it is entirely to this method's *credit* that we can recognize its limitations simply by applying itself to the question of which method to trust.¹¹

Modest Moral Methods

In this way, far from being suspicious of all modest evaluative methods, we should often welcome the result that an evaluative method is modest. And this is a point of real philosophical import. For there is good reason to welcome the result that the evaluative methods associated with some of the most important classes of practical reasons are modest in just this sense – including, most notably, the evaluative method associated with the point of view of morality itself.¹²

Why should we hope that the moral point of view is characterized by a modest evaluative method? The simple answer to this is that specifically moral reasons seem to be only part of the full story about what we ought to do, all things considered. Indeed, while moral reasons are surely very important, they do not *even* seem to *always* be the most important part of this story.

For example, many of us believe that we are sometimes rationally permitted to allow personal concerns to override moral ones in all-things-considered deliberation. The existence of such “agent-centered options”, even in cases in that involve non-trivial moral considerations, indicates that moral considerations are at most part of the full story about our reasons for action. And, more importantly, they also indicate that such reasons are not even the dominant element in this story in some cases – even when the moral reasons in question are not unweighty.¹³

Thus, there is a great deal of *prima facie* plausibility to the idea that morality is only one of a number of sources of legitimate reasons for action. As Thomas Nagel puts it:

¹¹ Of course, these remarks only apply if the methods in question do not display a cyclic form of modesty on which, say, M recommends N and N recommends M. Fortunately, there is relatively little reason to think that the cases under discussion below display this sort of cyclicalilty. Thanks to Jack Woods for pressing me to be explicit about this point.

¹² As will become clear, the structural point I want to make here applies to a wide variety of moral theories. So I'm not going to say very much at this moment about what is distinctive of moral reasons as such – since different accounts of morality will give very different answers to this question. Nonetheless, the interest of much of what follows is dependent upon our ability to make sense of the class of moral reasons in a manner that gives us good reason to care about such reasons and to treat them as a *distinct* class. Needless to say, many moral philosophers have thought of moral reasons in a manner that has these features. But the more we think of the space of practical reasons as lacking any deep distinction between moral and non-moral reasons, the less interesting these reflections will become. As will become plain I have some sympathy with such views. Still, even if we ultimately adopt of view of practical reasons on which there is not a deep distinction between moral and non-moral considerations, the idea that a plausible conception of the moral point of view will be modest about its own importance can help to undermine more moralistic conceptions of the relationship between morality and practical reasons in general. So even for someone with these leanings, these reflections may have considerable dialectical interest.

¹³ For an argument that any theory that accepts moral agent-centered options (or supererogation) must also accept that morality is modest about the weight of moral reasons, see Portmore's “Are Moral Reasons Morally Overriding?” Portmore presents this result as posing a *prima facie* challenge to such theories. Part of my goal here is to suggest that such a result is both more natural and more attractive than this might suggest – which I believe is also Portmore's final view of these issues. For a different argument in a similar, but more radical, direction see Harman's “Understanding Some Cases of Altruism as Permissible Moral Mistakes”.

Human beings are subject to moral and other motivational claims of very different kinds. This is because they are complex creatures who can view the world from many perspectives – individual, relational, impersonal, ideal, etc – and each perspective presents a different set of claims.¹⁴

If this is right, then there will be something misleading in any view that is too quick to single out moral reasons as providing us with anything like the whole story about what we ought to do all things considered. Thus, in investigating the class of moral reasons, we should not expect these reasons to provide anything like the whole story about rational deliberation. And as a result, we should not be overly concerned by the result that the evaluative method that formalizes what the moral reasons best support is modest about its own significance.

Indeed, if we are inclined to view morality in this way, we may feel that, “Like reason, the moral standpoint should try to recognize and explain its own limits.”¹⁵ If so, it is entirely to morality’s credit if the evaluative method that is distinctive of the moral point of view is modest in the manner just described. For, far from indicating a fatal inconsistency in the moral standpoint, this may simply indicate that morality itself is capable of recognizing its own limitations. For example, on such a view, morality itself would be capable of recognizing the existence of agent-centered options, which allow non-moral concerns to overrule specifically moral ones under some circumstances. Thus, far from being a reason to be suspicious of the moral point of view, its modesty might recommend this point of view to us all the more. For it might indicate that the moral point of view is itself sufficiently objective to appreciate the limited place of specifically moral reasons with the space of practical reasons more generally.

The Structure of Modest Moral Theories

If this is right, then many of us should welcome the result that the evaluative method that is distinctive of moral evaluation is modest in this way. For this may simply represent morality’s recognition of phenomena like the existence of agent-centered options – a recognition that is entirely to the credit of the moral point of view.

To understand the significance of this idea, it is important to note that if morality is modest, it will exhibit a fairly complicated structure. In particular, given any theory that is modest about the weight of moral reasons in this way, we can identify *two* evaluative methods (or better, a method and an über-method) that are associated with the moral point of view. First, there is what we may call the “method of moral evaluation” – namely, the evaluative method (i.e. M) that captures the immediate verdicts of the moral point of view itself. And, second, there is what we might call the “morally-approved method(s)” – namely, the evaluative method(s) (or more accurately, the über-method) that captures M’s recommendations about which method to trust in all-things-considered deliberation in this or that situation.¹⁶

Similarly, on such a view, there will be *two* classes of reasons that are associated with the moral point

¹⁴ “The Fragmentation of Value.” in *Mortal Questions*. (Cambridge: Cambridge University Press, 1979), 134.

¹⁵ *The View From Nowhere*. (Oxford: Oxford University Press, 1986), 203.

¹⁶ Note that this method may be somewhat disjunctive in character if M recommends a variety of different methods for different individuals in different circumstances.

of view. First there will be the reasons to which the method of moral evaluation M is itself responsive – i.e. the moral reasons in a *strict* sense. And, second, there will be what we might call the “morally-relevant reasons” – that is, the reasons that are given some weight by the morally-approved method(s).¹⁷ If morality is modest about the weight of moral reasons, the extent and weighting of the strictly moral reasons will not be identical with the extent and weighting of the morally-relevant reasons. Rather, there will be cases in which non-moral, but morally-relevant, reasons should trump moral reasons in rational deliberation, even morality’s own lights.

Not surprisingly, a similar potential ambiguity arises on such views with respect to talk of *moral requirements*. But it is most natural, given such a view, to say that we are morally required to do something only if the *morally-relevant reasons* decisively support this action.¹⁸ After all, in cases in which the (strictly) moral reasons decisively support some action, but the morally-relevant reasons do not, it seems misleading to speak of a genuine moral *requirement* to perform that action.

Thus, if we adopt this terminology, we can think of a modest moral theory as involving the following two elements:

- (i) A theory of the *narrowly moral reasons* (which is formalized by the method of moral evaluation).
- (ii) A theory of what is morally permitted and required that gives weight to all the *morally-relevant reasons* (which can be formalized by the über-method that generates the morally-approved methods).¹⁹

Thus, one important consequence of such views is that we cannot simply read moral requirements off of the moral reasons in isolation from the non-moral, but morally-relevant, reasons. For there may be cases in which quite weighty moral reasons are balanced out or even overridden from the perspective of the morally-approved method by even more significant non-moral reasons of this sort.

A number of features of ordinary moral thought speak to the idea that morality has this sort of general structure. For example, as Portmore and Harman argue, the difference between the narrowly moral reasons and the morally-relevant reasons that are recognized by the morally-approved method can be seen to underlie and explain the pervasiveness of supererogation as a moral phenomenon. And, as we have already noted, much the same is true of the idea that morality itself recognizes the existence of genuine agent-centered options. If morality is modest about the weight of moral reasons, there is nothing terribly surprising about either of these results. For if the moral requirements are not determined directly by the moral reasons alone, there is no reason to expect that we will always be morally required to do what the moral reasons best support.

More generally, any moral theory that is modest about the weight of moral reasons will exhibit a form of “indirectness”. In particular, on any such theory, what we are morally required to do will not be determined directly by the strictly moral reasons operating on their own. Rather, such

¹⁷ Here I draw on Portmore’s helpful terminology.

¹⁸ Note that this is a necessary and not a sufficient condition.

¹⁹ Note that, if morality in this guise is properly responsive to *all* the practical reasons, this über-method will represent the point of view of practical rationality itself. The important point here is that on views like those under discussion, it will also represent the verdict of the moral point of view on these questions as well. I’ll return to the relationship between the moral point of view and the point of view of practical reason below.

requirements will always be the product of both the moral reasons and the morally-relevant reasons operating together in the context of the morally-approved methods for all things considered deliberation. This sort of “indirect” structure is often thought to be a fatal flaw of moral theories like indirect consequentialism.²⁰ For, as Thomas Scanlon noted long ago, such views can easily seem to be an unstable compromise between conflicting philosophical impulses.²¹ But if the present reflections are on target, there is in fact reason to think that *any* satisfactory moral theory will have something like this “indirect” structure – reasons that are rooted in the deep structure of ordinary moral thought. If so, then *some* of the classical objections to indirect forms of consequentialism should, I think, lose much of their bite.²² And much the same will be true of some of the classical objections to deontological and virtue ethical views that have a similarly “indirect” structure.²³

Some Modest Moral Theories

Given this, it is not surprising that it is easy to find attractive moral theories that generate modest evaluative methods. For example, consider a form of Kantianism on which the method of moral evaluation (call this M) is concerned with the degree to which some alternative is justifiable to all rational beings. Notably, we can use M to evaluate *both* particular actions and general evaluative methods. That is, we can ask which particular actions would be best justifiable to everyone. And we can ask which evaluative methods it would be most justifiable for us to trust in general.

Pretty clearly, the answer to the second of these questions will consistently be something other than M itself. After all, the evaluative method that best passes a test of equal justifiability will presumably be one on which we all have some room to act in ways that do not maximize the degree to which our *particular* actions are justifiable to everyone.²⁴ Similarly, being consistently guided directly by M in one’s choice of actions appears to be incompatible with the sort of special concern that we ought to have for our friends and relatives.²⁵ Thus, it will be difficult to justify to those we are close to a general policy trusting M in our choice of actions. In short, the cases in which M recommends that we trust M (without any modification) seem to be relatively few and far between.

Something similar is also true of many attractive forms of consequentialism. For example, suppose we think of the method of moral evaluation M as concerned with the overall consequences of the possibility under evaluation. Then, much as in the Kantian case, we can use this evaluative method to evaluate both particular actions and more general evaluative methods. That is, we can ask which particular actions would generate the best consequences. And we can ask whether trusting some evaluative method in general would do so.

²⁰ For an excellent discussion of these issues see Brad Hooker. *Ideal Code, Real World: A Rule-consequentialist Theory of Morality*. (Oxford: Oxford University Press, 2000). See also Peter Railton's classic article, "Alienation, consequentialism, and the demands of morality," *Philosophy and Public Affairs* 13 (2):134-171 (1984).

²¹ “Contractualism and Utilitarianism” in Amartya Sen and Bernard Williams, eds., *Utilitarianism and Beyond* (Cambridge: Cambridge University Press, 1982): 103-28.

²² Of course, this only speaks to objections to the idea that moral theory should have this general *structure*. There are also particular objections to (say) indirect forms of *consequentialism* that it does not touch.

²³ Compare, for example, Samuel Scheffler’s discussion of the paradoxical character of agent-centered restrictions in *The Rejection of Consequentialism*. (Oxford: Oxford Clarendon Press, 1982).

²⁴ I.e. this method will allow for a sort of agent-centered options.

²⁵ See Bernard Williams. *Moral Luck*. (Cambridge: Cambridge University Press, 1981).

Once again, M is unlikely to be the method that scores best by the latter standard in almost any normal situation. For there seem to be decisive consequentialist grounds not to be the sort of person whose reasoning is moved solely by consequentialist considerations. After all, just as in the Kantian case, this would seem to render one incapable of many of the most valuable forms of human interaction.²⁶ And it also seems likely that we will maximize the good consequences of trusting a method by trusting a method that allows everyone some room for acting in ways other than performing the particular action that would bring about the best overall consequences. Thus, once again, we arrive at the result that M will recommend that we trust a method other than itself under normal conditions.²⁷

Modesty and Self-Effacement

As noted above, and as this discussion of consequentialism should make vivid, the fact that the evaluative method associated with a moral theory is modest in this sense is often taken to mean that this theory is “self-effacing”.²⁸ This has often been taken to be a serious disadvantage of theories of this sort – at least insofar as this propensity towards self-effacement becomes extremely widespread.²⁹ But, once again, this is not the only way of interpreting the fact that the evaluative method that is distinctive of the moral point of view is modest in the sense defined above.

Morality is self-effacing just in case it recommends that we believe something other than the truth about it. Or, to put things more broadly, a moral theory is self-effacing insofar as we ought, by its own lights, not to accept it. This question – like the questions I am interested in – concerns a

²⁶ See Michael Stocker. “The Schizophrenia of Modern Ethical Theories,” in *Virtue Ethics*, eds. Roger Crisp and Michael Slote (Oxford: Oxford University Press, 2002): 66-78.

²⁷ Similar results follow from a sentimentalist characterization of the moral point of view like Hume’s – provided that the sentimental responses that characterize the moral point of view do not respond most positively to individuals whose choice of actions is dominated by those very sentiments. According to at least Hume this is the case – for being responsive to one’s moral sense plays a secondary role in Hume’s own conception of virtue. And this will be true on most views that characterize the moral point of view in terms of a sentimentally constituted moral sense. After all, much as Hume lays out, while our moral emotions and sentiments may approve to some degree of individuals whose actions are determined exclusively by their moral emotions and sentiments, they will generally approve much more of individuals whose actions are determined by some mix of their moral sentiments and motivations of a non-moral sort, so that they sometimes act in ways contrary to the recommendations of their moral sentiments alone. As Hume noted, someone with a healthy mixture of private, public, and moral concerns receives the approval of our moral sentiments much more than a “morally pure” Cato who has completely subjugated his person life to his understanding of virtue. On this empirical point, it seems to me that Hume was surely correct. And this feature of our moral sentiments seems unlikely to disappear with the sorts of idealization that sentimentalists account of the moral point of view generally involve.

²⁸ See, for example, the classic discussion in Derek Parfit’s *Reasons and Persons* (Oxford: Oxford University Press, 1984).

²⁹ How widespread the self-effacement is is crucial here, since it is plausible that all attractive moral theories will be self-effacing under *certain* circumstances. (C.f. Simon Keller, “Virtue Ethics is Self-Effacing.” *Australasian Journal of Philosophy* 85:2 (2007): 221-232; and Ben Eggleston, “Rejecting the Publicity Condition: The Inevitability of Esoteric Morality.” *Philosophical Quarterly* 63:250 (2013): 29–57.) What is potentially problematic, I take it, is not this result, but rather the result that a moral theory is self-effacing in any normal situation. Or, in other words, what seems problematic is the result that it would always be better morally speaking for us to believe something other than the truth about morality.

morality's verdict about itself. But my focus is on something else, namely morality's verdict – not about whether we should accept it as the truth about morality – but rather about whether moral reasons are overriding in all things considered deliberation. And morality can easily recommend that we regard moral reasons as less than overriding, without recommending that we have false views about such reasons.

For example, consider a moral theory T, on which the method of moral evaluation M is modest in the sense defined above. As just noted, this might simply mean that T is self-effacing in the standard sense. That is, focusing on my current situation for the sake of simplicity, it might mean:

$M(\text{I believe } T) < M(\text{I believe } R)$, for some moral theory $R \neq T$.

But this is not the only possibility here. For M might instead recommend that I continue to believe T - and, indeed, that this belief continue to be motivationally efficacious – even while recommending that I trust some other method in all-things-considered deliberation, so that:

$M(\text{I believe } T \text{ and this belief is motivationally efficacious}) > M(\text{I believe } R)$, for all moral theories R.

Perhaps surprisingly, there is nothing incoherent about this second possibility. For if our moral theory does not take moral reasons to be consistently overriding, there is no conflict between having a motivationally efficacious belief in this theory and being guided in one's all things considered deliberations by a principle other than “do what the moral reasons best support”. Thus, the fact that M recommends that I trust another evaluative method in such deliberation need not imply that T is self-effacing.

Indeed, the hypothesis that morality is modest about the weight of moral reasons generally seems a more charitable interpretation of these sorts of formal results than the idea that morality is systematically self-effacing. For example, consider Michael Stocker's famous argument that consequentialism is objectionably self-effacing.³⁰ As alluded to above, this argument rests on the plausible claim that there are decisive consequentialist grounds not to be the sort of person who is motivated only by consequentialist considerations. In the literature this is often taken to show that consequentialist moral theories recommend their own rejection. But this follows only if we assume that the theory in question takes moral reasons to be consistently overriding. The alternative, of course, is to view consequentialism merely as a theory of moral reasons – and to extend this theory so that it also acknowledges that such reasons can sometimes be overridden by reasons of other sorts in rational deliberation. If viewed in this way, consequentialism might recommend that moral reasons sometimes be overridden by non-moral considerations, without thereby recommending that one reject it *qua* the true moral theory. Thus, we can react to arguments like Stocker's in at least two ways: either by treating morality as *both* overriding and self-effacing or by treating it as *neither* self-effacing nor overriding.

In many cases, the latter reaction seems to me far more palatable. For example, consider a form of indirect consequentialism (I) that has the following structure. First, there is a standard consequentialist account of the nature of specifically moral evaluation – yielding a method of

³⁰ “The Schizophrenia of Modern Ethical Theories.”

evaluation M. But second, this method is not taken by T to directly determine the weight that moral and non-moral considerations should have in all-things-considered deliberation. Rather, T claims that these weights should be determined by the consequences associated with *giving these reasons a certain weight in deliberation* – so that the T-approved method for all-things-considered deliberation is just the method, the trusting of which by us would generate the best results according to M. Such a theory would avoid the standard worries about self-effacement in just the manner sketched above. For while such a view would deliver the result that:

There is an evaluative method $N \neq M$, such that $M(I \text{ trust } M) < M(I \text{ trust } N)$.

It would not follow from this that:

$M(I \text{ believe } T) < M(I \text{ believe } R)$, for some moral theory $R \neq T$.

Rather, in deliberating in accordance with the non-consequentialist method N, we would be doing exactly what T *qua* indirect consequentialist theory asks us to. Thus, one way for consequentialists to avoid worries about self-effacement is to embrace the idea that morality is modest about the weight of moral reasons.³¹

Modesty and the Weight of Moral Reasons

This is an important possibility. But what interests me most about these ideas at present are their implications for debates about whether moral reasons have overriding strength in all things considered deliberation. This idea is often objected to on grounds that lie outside the scope of morality itself. So, for example, Susan Wolf famously attacked the claim that our lives should be dominated solely by moral considerations by pointing to the deeply unattractive character of the lives of such “moral saints” from the perspective of personal perfection.³² But such “external” attacks on the overridingness of moral considerations have obvious limitations. For while such complaints will appeal to those of us who are ready to view morality in this way, to a defender of the overridingness of morality, they may seem to miss the point. Indeed, someone with such “moralistic” tendencies is likely to retort that, if Wolf is right, then so much the worse for the standpoint of personal perfection.

In order to do better than this, we need to find a way of objecting to the idea that moral considerations are overriding from within the standpoint of morality *itself*. Attempts to do just this are not absent from the philosophical literature. Perhaps the most notable of these is Thomas Nagel’s argument for the idea, cited above, that “the moral standpoint should try to recognize and explain its own limits.” Nagel does offer a *moral* argument for this claim. But for Nagel the crucial issue here is our recognition, from a moral point of view, of our own limited capacity to be moved

³¹ Compare and contrast this sort of view with the views that Jonathan Dancy attacks in *Moral Reasons* (Oxford: Wiley-Blackwell, 1993), chapter 13, which attempt to draw a distinction between two classes of *moral* value or reasons.

³² For defenses of versions of this idea see Susan Wolf’s “Moral Saints” in *Journal of Philosophy* Vol. 79, No. 8 (1982), 419-439. Compare Philippa Foot’s “Are Moral Considerations Overriding?” in her *Virtues and Vices* (Oxford: Oxford University Press, 2003) and Samuel Scheffler’s *Human Morality* (Oxford: Oxford University Press, 1992). Finally, see of course Sidgwick’s discussion of the “dualism of practical reason” in *The Methods of Ethics* (1907).

by purely moral considerations. Given this, Nagel argues that insofar as morality conceives of itself as a system of principles the point of which is to guide human action, it will recognize the need to moderate its demands insofar as this is required for it to fulfill this role.³³ Thus, morality itself, according to Nagel, will recognize the need to strike “a bargain between our higher and lower selves” with respect to the manner in which it regards its own distinctive concerns.

This argument, and other arguments like it, does present us with a sort of moral argument against regarding moral considerations as overriding. But it is a rather backhanded argument for this conclusion. For the most it supports is the claim that we must *regrettably* adjust the demands of morality to make them motivationally efficacious for human beings “as we find them”.³⁴ It is at this point that the potential modesty of the moral point of view becomes relevant. For if the moral point of view is modest about the weight of moral reasons, then the idea that moral reasons consistently trump non-moral considerations does not withstand scrutiny *even* from the moral point of view itself. And this is in no way regrettable, even from the moral point of view. Rather, from the start, we will find the dominance of narrowly moral considerations and values unattractive, even when we view this question from a moral perspective. Thus, it is not merely that morality recognizes that it must strike “a bargain between our higher and lower selves” of the sort Nagel has in mind. Rather, it approves of a balance between these two “selves” prior to any “bargaining” of this sort. In short, on such views, Wolf’s “moral saints” aren’t even *moral* saints. For, at least as these theories see things, a thoroughgoing adherence to a strictly moral perspective on the world is contrary to morality’s own conception of its practical importance.

In this way, the possibility that morality is modest about the weight of moral reasons seems to provide a much more secure basis for objecting to the overridingness of moral reasons than arguments like Nagel’s. For given such a conception of morality, we can argue against the overridingness of moral reasons as follows:

1. Moral reasons are consistently overriding only if moral reasons seem to be consistently overriding from the moral point of view.
2. Moral reasons do not seem to be consistently overriding from the moral point of view.
3. Therefore, it is not the case that moral reasons are consistently overriding.

Or alternatively:

1. Moral reasons are consistently overriding only if morality itself takes moral reasons to be consistently overriding.
2. Morality does not take moral reasons to be consistently overriding.
3. Therefore, it is not the case that moral reasons are consistently overriding.

Or finally, making use of the reading of “moral requirements” presented above:

1. Moral reasons are consistently overriding only if morality itself takes moral reasons to be consistently overriding.

³³ Similar considerations play a role in Samuel Scheffler’s argument for the thesis that morality is “moderate” in his *Human Morality*. Plainly my thoughts about these matters are heavily indebted to the work of both Nagel and Scheffler on these topics.

³⁴ It is here, I think, that arguments like Wolf’s are more compelling.

2. Morality takes moral reasons to be consistently overriding only if moral requirements are purely a function of the moral reasons.
3. Moral requirements are not purely a function of the moral reasons.³⁵
4. Thus, morality does not take moral reasons to be consistently overriding.
3. Therefore, it is not the case that moral reasons are consistently overriding

Of course, as noted before, for a moral theory that is modest about the weight of moral reasons, there are in fact two classes of reasons that may be associated with the moral point of view: (i) the moral reasons in a strict sense, to which the method of moral evaluation is itself responsive, and (ii) the morally-relevant reasons, which are given some weight by the morally-approved method. The arguments just given speak against the overridingness of moral reasons in the first, strict sense. But it is important to note that they do not speak against the overridingness of the morally-relevant reasons. For, of course, there is nothing in the moral point of view (even on such accounts) that speaks against the idea that these are all the considerations that are relevant to what one ought to do all-things-considered.

A Modest Morality as the Foundation of Practical Rationality?

Of course, it might be the case that the morally-approved method is modest about whether we should trust it in just the same way the method of moral evaluation is.³⁶ But putting this possibility aside, there is nothing in such views to conflict with the claim that the morally-approved method *is* the whole story about matters of practical significance.³⁷ Thus, even on such a view, it remains possible to view morality as the *final* or most fundamental court of appeals on matters of practical reasoning. For it remains open to us to take morality's verdicts about which method we should use in deliberation - and so, indirectly, about the relative weight of moral and non-moral reasons - to be final and authoritative. And if we do so, then while morality would be modest about the weight of moral reasons, it would not be modest about its own significance in a *deeper* sense. For it would regard itself as determining the weightiness of moral and non-moral reasons alike (via determining the correct deliberative method to use *all things considered*). Thus, it is perfectly possible to maintain both that the moral point of view is foundational in this way and to maintain that it is modest about the significance of narrowly moral reasons.

Thus, there is no reason why someone who believes that morality is modest about the weight of moral reasons needs to reject the idea that the morality *explains* the weight that different considerations have in practical reasoning. In this way, morality can be modest in the sense at issue here and still be said to be the "supreme principle of practical reason" in one important sense of this phrase. Indeed, once we recognize this, there is good reason to read even someone as seemingly moralistic as Kant as endorsing a version of the thesis that morality is modest in the sense of interest to us here. For while Kant surely does regard the moral law as providing the basis on which all practical reasons in some sense rest – the manner in which it carries out this task is by no means direct or immediate. Rather, as Barbara Herman notes in her excellent discussion of these issues, for Kant:

³⁵ Here again my discussion connects up with Portmore's and Harman's more detailed discussion of this idea.

³⁶ Although here again the danger of cyclic forms of modesty appears.

³⁷ But see Andrew Sepielli's "What to Do When You Don't Know What to Do When You Don't Know What to Do." forthcoming in *Nous* for the very real possibility that this sort of modesty goes very deep indeed.

In moving elements of own happiness into the space of moral[ly-relevant] reasons for each agent, happiness is not subsumed by morality, as if the determination of what makes my life go best is to be made impersonally. What is absorbed into morality is the status of the pursuit of happiness.³⁸

In this way, we can read Kant as claiming that the foundational role of morality is a matter of its ability to ground the class of morally-relevant reasons – and not the, much more stringent, claim that only *narrowly* moral reasons are genuine reasons for action.

Nonetheless, even on views such as Kant's, the idea that morality is modest about its own practical significance provides us with the materials for an *internal* critique of the idea that moral reasons are always overriding. And this should be attractive to us, whether or not we are inclined to treat morality as having the sort of foundational role just described. For it provides us with a means of engaging with those who advocate for a more stringent conception of morality on something like their own terms – something that should be welcomed by all friends of a more humane perspective on practical matters, whatever their own opinion of the place of morality within the practical domain.

Of course, this will only be true if we ought to accept a view of morality on which the moral point of view is modest in this sense defined above. Obviously, there is no way to answer this question without taking a stand on controversial issues about the nature of morality and its place in human life. But such a view of morality seems to me to have much to recommend it over an alternative on which morality regards moral reasons as consistently overriding non-moral ones.³⁹ And, as just noted, we can accept a modest conception of morality, while also taking morality to explain why it is that certain non-moral considerations have the weight they do. Thus, taking morality to be modest in this sense does not require that morality abdicate her throne, it only requires that her rule over the practical sphere be somewhat indirect than we might have thought.

Moreover, with this simple step, we can avoid the moral views that Hume famously condemned as the product of “gloomy, hair-brained enthusiasts”. Of course, the gloomy enthusiasts Hume had in mind were the product of extreme tendencies in religion – whereas today there are as likely as not to be the product of extremes of secular philosophical argument. But in either case, like Hume, I believe that we should be suspicious of their claims on *moral* grounds. For if I am right, there is good reason to that morality is modest about its own significance. And if it is, then the best way to do justice to the authority of morality is to give it a foundational role in our theory of practical reasoning – and *not* to make its immediate demands overriding in each and every case.

³⁸ “The Scope of Moral Requirement” (*Philosophy and Public Affairs*, 2001), 244. Note that Herman calls *all* of the reasons that are informed in this by morality as “moral”, while I have used the label “morally-relevant” to refer to this class. But this difference between us seems to me to mainly be terminological in character.

³⁹ See again the arguments in Portmore and Harman for further support of these claims.