

ACCEPTED VERSION

Garrett Cullity

Virtue ethics, theory, and warrant

Ethical Theory and Moral Practice, 1999; 2(3):277-294

DOI: [10.1023/A:1009903128864](https://doi.org/10.1023/A:1009903128864)

The final publication is available at link.springer.com

© 1999 Kluwer Academic Publishers

PERMISSIONS

<http://www.springer.com/gp/open-access/authors-rights/self-archiving-policy/2124>

"Authors may self-archive the author's accepted manuscript of their articles on their own websites. Authors may also deposit this version of the article in any repository, provided it is only made publicly available 12 months after official publication or later. He/ she may not use the publisher's version (the final article), which is posted on SpringerLink and other Springer websites, for the purpose of self-archiving or deposit. Furthermore, the author may only post his/her version provided acknowledgement is given to the original source of publication and a link is inserted to the published article on Springer's website. The link must be accompanied by the following text: "The final publication is available at link.springer.com".

14th May, 2015

<http://hdl.handle.net/2440/90971>

VIRTUE ETHICS, THEORY, AND WARRANT*

Garrett Cullity

ABSTRACT

Are there good grounds for thinking that the moral values of action are to be derived from those of character? This "virtue ethical" claim is sometimes thought of as a kind of normative ethical theory; sometimes as form of opposition to any such theory. However, the best case to be made for it supports neither of these claims. Rather, it leads us to a distinctive view in moral epistemology: the view that my warrant for a particular moral judgement derives from my warrant for believing that I am a good moral judge. This view seems to confront a regress-problem. For the belief that I am a good moral judge is itself a particular moral judgement. So it seems that, on this view, I need to derive my warrant for believing that I am a good moral judge from my warrant for believing that I am a good judge of moral judges; and so on. I show how this worry can be met, and trace the implications of the resulting view for warranted moral judgement.

KEYWORDS: Virtue ethics, Ethical theory, Moral epistemology, Moral judgement

* Earlier versions of this material were presented to audiences at the Universities of Sheffield and St. Andrews, as well as the 1998 British Society for Ethical Theory conference. For comments that led to rethinking, thanks to Peter Carruthers, Chris Hookway, David Bell, Ralph Wedgwood, Nick Zangwill, Philip Stratton-Lake, Rosalind Hursthouse, and Matthew Kieran. For especially helpful written comments, thanks to Mark Nelson, Brad Hooker, Jimmy Lenman and Andrew Moore.

I: INTRODUCTION

One claim often made distinctive of “virtue ethics” is this:

(V) The moral values of action are to be derived from those of character.

This can be read different ways, since saying that one value is “to be derived” from another can mean different things. On a first reading, (V) asserts a relationship of constitution between the value-tokens in question: what *makes* a given action have the value it has is the value of a related state of character. On a second, by contrast, it claims a relationship of justificational warrant between judgements attributing those values. These two relationships — call them “constitutive justification” and “warrant for judgement”, respectively — are clearly distinct: I am not always warranted in judging truly, and am sometimes warranted in judging falsely, so what makes something the correct thing to think and what warrants me in thinking what I am warranted in thinking must be different.

This essay examines whether there is a plausible case for (V), and its relation to ethical theory. (V) is often presented as a distinctive kind of normative ethical theory, to be contrasted with the consequentialist and deontological alternatives. Instead of assessing the rightness of actions in terms of their consequences or the act-types to which they belong, a “virtue theory” assesses them in terms of the states of character from which they issue.¹ However, incompatibly with this, the emphasis on the priority of character-over act-assessment is sometimes associated with a repudiation of ethical theory altogether.² Is either of these views plausible?

¹ See e.g. Michael Slote, “Virtue Ethics,” in Marcia W. Baron, Philip Pettit, and Michael Slote, *Three Methods of Ethics* (Oxford: Blackwell, 1997); Rosalind Hursthouse, “Normative Virtue Ethics,” in Roger Crisp (ed.), *How Should One Live?* (Oxford: Clarendon Press, 1996).

² For works which suggest this latter position to varying degrees, see Bernard Williams, *Ethics and the Limits of Philosophy* (London: Collins, 1985), Ch.10; Annette Baier, *Postures of the Mind* (Minneapolis:

I begin with the second. Is there a tenable position linking (V) to the rejection of ethical theory? We need to identify the target of the attack to assess its success. I take it that its essential features are captured by the following:

An ethical theory is a morally substantive set of propositions such that:

- (a) It is claimed that those moral judgements that follow from it are correct.
- (b) It is highly systematic, relative to the judgements for which it supplies a theory.
- (c) It is highly unified, relative to the judgements for which it supplies a theory.

This makes the issue of whether a set of propositions constitutes an ethical theory a matter of degree. A structure is *systematic*, I shall assume, to the extent that its members are linked by inferential chains. A and B are linked by an inferential chain when they are the first and last members of a sequence each member of which bears a direct inferential relationship to the previous one – either supporting or being supported by it. A structure is *unified* to the extent that its members have common inferential sources.

The inferential relationships claimed by an ethical theory are relationships of constitutive justification. But a common rationale for seeking such a theory concerns warrant for judgement. If we can be warranted in advocating an ethical theory, we are warranted in making the judgements that follow from it; so it can seem promising to look to an ethical theory for a source of warrant for contentious moral judgements. This can even seem the only such source — without a warrant for advocating such a theory, the aim of warranting many contested judgements must be abandoned, since they will be no better supported than their contraries. The virtue-ethical attack on ethical theory, as I shall characterize it, maintains that this idea rests on a confusion, and, in particular, on a misguided picture of the role of rules in the justification of moral judgement.

University of Minnesota Press, 1985), pp. 209–20, 234-6; Edmund L. Pincoffs, *Quandaries and Virtues: Against Reductivism in Ethics* (Lawrence, Kansas: University Press of Kansas, 1986); and, especially, John McDowell, “Virtue and Reason,” *Monist* 63 (1979), Sect.4.

My first conclusion is the rejection of this claim. In any sense in which (V) is plausible, it leaves plenty of scope for theorizing in ethics. But my second is that it does not make sense to think of such an activity as furnishing a “virtue theory”. In doing so, I oppose the other line concerning ethical theory taken by proponents of (V).

These conclusions are entirely negative. But there is a more positive side to the essay. For despite these conclusions, a strong case can be made for (V) itself, at least on one interpretation — an interpretation, moreover, on which it supplies a plausible framework for moral epistemology.

II: CHARACTER, RULES, AND THEORY

Just how are the values of actions to be derived from those of character, as proposed by (V)? In particular, from *whose* character are we to derive them?

(V) has no plausibility when read as the claim that the moral value of an action is to be derived from that of the character of the agent performing it — for actions can be performed “out of character”. To make (V) credible, it must be interpreted along these lines:

(V1) The rightness of an action in a given set of circumstances is to be derived from its being one that a fully informed and virtuous agent would perform in those circumstances (if acting in character).

Now this may seem to run into a familiar kind of problem: one common to all attempts to explain norms governing actual agents in terms of descriptions of ideal ones.³ For surely the appropriate action for *me*, a less than fully virtuous agent, differs in various

³ The now-standard response to this problem is to talk of the advice that ideal agents would give actual ones: see e.g. Michael Smith, *The Moral Problem* (Oxford: Blackwell, 1994), pp.150–51ff. My own response in the text amounts to a version of this. For the difficulties with offering this kind of claim as an *analysis* of rightness see Simon Blackburn, *Ruling Passions* (Oxford: Clarendon Press, 1998), pp.104-19.

ways from what a perfectly virtuous agent would do. I ought to make amends to people I've let down. I ought sometimes to do things to help myself become more virtuous. And my actions ought to take account of my own moral imperfection: given my own limited moral courage, for example, it might be irresponsible to take on a role in which others importantly rely on me to stand up to pressure. But a perfectly virtuous agent would never need to do these things.

These would be serious problems if (V1) referred us, in talking of the action “that a virtuous agent would perform”, to an agent who was exceptionlessly virtuous throughout his life. But (V1) should not be read that way. Think instead of an agent with my past, and who knows he will revert to my psychological make-up once his decision is taken. *These*, after all, are my circumstances: the circumstances of my past and future moral imperfection. But now the examples just given cease to be problematic: the rightness-claims they contain are all endorsed by (V1). The last of them does raise a further point. Should I commit myself to a role demanding greater moral courage than I think I possess? Declining is the right thing *for me* to do. However, there also remains a sense in which the morally courageous action is *the* right thing to do — it's just that I am not the right person to do it. (V1) usefully spans both answers. The first corresponds to the action of a virtuous agent who knows he will revert to my psychological make-up; the second, one who knows he will not.

(V1) is not simply refuted by a familiar line of objection. But how can it be supported, and what is its relation to the attack on ethical theorizing? One influential argument answers these questions by attacking a picture of the role of rules in ethical justification.⁴

⁴ For the argument that follows, compare John McDowell, “Virtue and Reason,” *Monist* 63 (1979), Sect.4; also his “Non-Cognitivism and Rule-Following,” in Stephen H. Holtzman and Christopher M. Leich (eds), *Wittgenstein: To Follow a Rule* (London: Routledge & Kegan Paul, 1981). The problem discussed here should be distinguished from *the* “problem of rule-following” (much discussed in connection with

On this picture, justification presupposes the existence of rules. If the application of a moral concept C to a given object is to be justified, it must instantiate a rule for applying C, and if I am to be justified in applying C to that object, I must be following that rule. For only the existence of a rule governing a practice prevents it from being arbitrary, and the arbitrariness of a practice precludes any talk of justification in relation to it.

Now there is a way of taking this that makes it incontrovertible. That there is a right way and a wrong way of applying a given concept is just *equivalent* to the existence of a rule ruling in some ways and ruling out others. The principal claim here concerns constitutive justification:

(a) What makes it right to apply concept C to this object is its instantiating the rule for the application of C.

And from this, a claim concerning warrant for judgement seems to follow. Provided I possess the concept of following a rule, claim (a) tells me that I will have a warrant for believing that I am right to apply C to this object whenever I have a warrant for believing that there is a rule for the application of C that I would be following if I did so.⁵ If so, then, at least for those who do have that concept, we also have:

(b) What warrants me in judging that C applies to this object is my warrant for judging that I am following the rule for the application of C.

However, it can be tempting to construe (a) and (b) in a way which makes them far from incontrovertible. This happens when it is assumed that the rule they mention must be independently articulable — articulable independently, that is to say, of the rule:

(R) Identify instances falling under the concept C.

Wittgenstein) concerning what kind of entity a rule can be if it is both to determine its applications, and be accessible to us.

⁵ I could be warranted in thinking I was applying a concept correctly without being warranted in believing anything about rule-following, if I lacked the concepts necessary for forming the latter belief.

This assumption can tempt people to think that if the practice of applying C is to be non-arbitrary, there must be some *further* rule governing it, and that I am warranted in judging that I am rightly applying C only if I am warranted in thinking that I am guided by such a rule. And according to “virtue-ethical” opponents of ethical theory, it is this thought that encourages the view that warranting ethical judgements about contentious cases requires a theory. In order to warrant my judgement about the application of C in a contested case, I must produce and justify a rule supporting my practice; and such rules, and their justifications, are supplied by ethical theories.

The complaint is that this further reading of (a) and (b) is not only unsupported by the thoughts about arbitrariness; it looks highly dubious. If the application of a concept is not to be arbitrary, there must be right and wrong ways of applying it, so there must be a rule ruling in right and ruling out wrong ways of applying it. But all this requires is that, if there are to be right and wrong ways of applying C, then (R) must itself constitute a genuine rule, rather than ruling out nothing. It does not require a further, independently articulable rule for when one counts as following (R). After all, a general requirement of this kind would apply to the concepts employed in any further rule, producing either circularity or a vicious regress — vicious, because it would yield an infinite and therefore unfulfillable chain of relationships of dependence.

There must, therefore, be some concepts for which there are no independently articulable rules; and if so, we have been given no argument for thinking that moral concepts must be governed by such rules. From this conclusion, which concerns constitutive justification, a conclusion concerning warrant for judgement follows, as before. If no independently articulable rule governs the application of a concept, my being warranted in applying it as I do can hardly depend on my being warranted in judging that there is such a rule.

The upshot of this argument is that when I face a conflict between moral reasons — pro tanto reasons to assign incompatible moral values to an action — I should not look to a theory to provide me with a further, fully determinative rule for resolving the conflict,

by assigning relative strengths to those reasons. At the very least, the onus lies with a proponent of such fully determinative rules to produce a plausible example, once the thought that they *must* exist has been undermined. What I need instead, according to this line of argument, is moral sensitivity, good moral judgement: an appreciation of the true moral value of things. Having this does involve following a rule — but only the rule:

Give every moral reason its due moral importance.

This is what I must do in order to make the right judgement when my moral reasons conflict; but we have been given no case for believing that that requires me to be following some further rule.

How, on this view, can I be *warranted* in evaluating an action as I do? This seems to require a warrant for thinking that, when I make my evaluation, I am exercising good moral judgement. The difficulties with this view will be our next concern. But first, notice how it delivers a version of (V1).

It does so provided our usage of virtue-terms makes moral discernment a necessary condition of virtue. This is by no means the only usage: indiscriminating action is commonly described as virtuous but misguided — as misplaced kindness, say, or excessive loyalty. But (V1) cannot be using virtue-terms that way; for clearly, on that usage virtuous and fully informed action need not be right. (V1) must be presupposing a different usage, according to which you only count as possessing a moral virtue to the extent that your attitude towards the situations it covers is properly discriminate: if your devotion to a friend or cause is excessive, then it displays not the virtue of loyalty but the vice of partisanship. This usage need not make good judgement necessary for virtue: it can allow that, if my beliefs about a situation are reasonable but incorrect, that compromises my judgement but not my virtue.⁶ Virtue always involves discernment, on

⁶ Nor is it sufficient. A courageous person must be able to appreciate what is the right thing for an agent to do in situations of danger or discomfort. But this is not sufficient for courage. She must be able to see this when she is *in* such situations, and must not only see what is the right thing to do but must do it as well.

this usage, but discernment must be supplemented with full information to produce good judgement. That is, on the usage of virtue-terms (V1) presupposes, the actions a good moral judge would think right are those a fully informed and virtuous agent would perform (if acting in character). Presumably, not everyone is warranted in believing this. But for those of us who are, we have a warrant for believing that a fully informed and virtuous agent would perform an action whenever we have a warrant for believing that a good moral judge would think it right. The earlier argument concluded that what warrants me in judging that a given action is right is my warrant for thinking that I am exercising good moral judgement in so judging. So, putting these claims together, we arrive at:

(V2) What warrants me in judging that this action in these circumstances is right is my warrant for judging that a fully informed and virtuous agent would perform this action in these circumstances.

And this is claim (V1), interpreted as concerning warrant for judgement.

This should not, however, lead us to take (V1) seriously when taken as a claim about constitutive justification. Indeed, if (V1) presupposes a usage according to which an action only qualifies as expressing a virtue to the extent that it is not morally objectionable, it is hard to see how what *makes* an action right could be its expression of a virtue.⁷ On the contrary, this usage seems to be taking an action's rightness to be constitutively prior to its virtue-expression.

Accordingly, (V2) will be our focus in what follows. Can moral epistemology be satisfactorily pursued along these virtue-ethical lines? Let us investigate.

⁷ I say this is hard to see; but I think there is at least one case in which we can find such support. See my "Moral Character and the Iteration Problem," *Utilitas* 7 (1995).

III: WARRANTING VIRTUE-CLAIMS

How can I be warranted in thinking a given action has moral value M? On the virtue-ethical view on offer, what warrants me in thinking this is my warrant for thinking that good moral judges would think this.

Now one way of taking this makes it easily rejectable, and another makes it trivial. The first takes it as a claim concerning what it is to be warranted in making a moral judgement — what warrant for such a judgement consists in. But as an account of this kind, what has been offered is clearly hopeless. It would be viciously circular, in purporting to explain the constitution of warrant for thoughts with a moral content by appealing, in the explanans, to the notion of warrant for thoughts with a certain specific sort of moral content, concerning good moral judges. A second way of taking the claim is as the modest biconditional:

I am warranted in thinking that this action is M if and only if I am warranted in thinking that good moral judges would think that this action is M

But this is modest enough to be trivial. Good moral judges are simply those who make non-accidentally correct moral assessments. So this biconditional simply reduces to the claim: I am warranted in thinking that this action is M if and only if I am warranted in thinking that people who make non-accidentally correct moral assessments would think that this action is M.

But the virtue-ethicist's claim should be read in neither of these ways. It is not a meta-epistemological claim concerning the constitution of moral warrant, but a substantive position in moral epistemology, concerning the conditions under which particular moral judgements are in fact warranted.⁸ And as a substantive position of this kind, it goes beyond the trivial biconditional to make a claim of epistemological priority. When I am warranted in judging that this action is M, what warrants me in that judgement is my warrant for judging that good moral judges would think this. It is to my warrant for the

⁸ Still less is it a kind of ideal observer theory of the constitution or extension of moral value itself.

claim concerning good moral judges that I must look to supply a warrant for my own moral views. “The virtue-ethicist’s view”, in what follows, will refer to this claim of epistemological priority, as drawn from the rule-following argument.

The virtue-ethicist’s view makes the basic task for moral epistemology the warranting of judgements about what good moral judges would think. But how can I acquire a warrant for *these* judgements? It seems I must do three things:

- (a) Produce an account of the characteristics distinctive of good moral judges.
- (b) Justify that account.
- (c) Justify the claim that judges with those characteristics would make the judgement in question.

And now we seem to have a dilemma. Either the account at (a) amounts to a fully determinative rule of the form:

All and only those people possessing characteristics XYZ are good moral judges;

or it gives us a prima facie rule of the form:

X, Y, and Z provide presumptive reasons to think that someone is a good moral judge.

But fully determinative rules are precisely what the virtue-ethicist is questioning. Non-arbitrary standards for good moral judges do not presuppose a fully determinative rule, articulable independently of the notion of a good moral judge. This does not prove that there is no such rule, but it casts the onus on their proponents to produce a plausible candidate, and it is hard to be optimistic about the prospects of producing one that is both fully specific and exception-proof.

If so, account (a) can supply at most a set of prima facie rules for identifying good moral judges. But if it supplies only prima facie rules, it cannot on its own license conclusions concerning who is a good moral judge. In order to apply those rules correctly, I need good moral judgement. And that seems to produce a vicious regress. Account (a) was being sought to enable me to warrant my judgements concerning good moral judges.

But it turns out that I can only apply the account correctly if I myself am a good moral judge (of good moral judges). So I can only be warranted in thinking that I am drawing the right conclusions if I am warranted in thinking that I am a good moral judge. But what warrants me in thinking that? This was the question we wanted an answer to in the first place.

To see our way out of this dilemma, we should begin by trying to sketch an account of good moral judges.

IV: GOOD MORAL JUDGES

What are the distinguishing characteristics of a good moral judge? We can put the same question by asking about the characteristics of good moral judgement — provided we are clear that “judgement” here refers to the faculty of judging, or the exercise of that faculty on particular occasions, rather than the thing judged. (What I judge can happen to be correct even if I am a bad judge.)

Rather than attempting a fully argued answer, I shall offer a list of five features that at least provide a starting-point for an account of good moral judgement, and an explanation of the sources of support for it. The aim is simply to show the reasonableness of thinking that a proper account could be given. But we shall then see how this allows us to solve our dilemma.

The first feature is this: basing one’s judgement on full attention to all and only the morally relevant features of the object judged. If, in thinking about a situation, you have simply overlooked an issue of integrity, say, or gratitude that it involves, then that is a ground for faulting your judgement.⁹ The same applies if you have registered these issues, but only glancingly; or if you are influenced by a morally irrelevant issue, such as the appearances of the people involved. The relevance of this first characteristic to good

⁹ I.e. (as before), fault your judging, not what you judge.

moral judgement is obvious. To the extent that one is warranted in thinking a feature of an object morally relevant, one is warranted in faulting the judgement of someone who fails fully to attend to it, and to the extent that one is warranted in thinking a feature morally irrelevant, one is warranted in faulting the judgement of someone who gives it moral weight. What is less obvious is what to say about the epistemology of judgements concerning moral relevance: we shall need to return to that later.

Full attention to all and only morally relevant features does not suffice for good moral judgement, however. One must think about them in the right way, and assign them the right relative importance, in reaching a conclusion about the object of assessment. What are the characteristics of an agent who is also doing this? Clearly, she must be free of inferential error or inconsistency — this is a second feature of good moral judges. For a third and fourth, it seems we should add sympathy and impartiality. That someone's interests are affected by a proposed course of action is a morally relevant feature of it; if I fully attend to this, but lack sympathy, that suffices to vitiate my judgement about the action. And partiality to myself or those I love can readily distort my judgement about what is right, even if it is flawless in the other respects.

A fifth feature can be added to our list, provided it is clear that it is not a *constituent* of good moral judgement, but only an indicator of it: this is a person's relevant moral experience. My failure to appreciate the right way of taking the various morally relevant features of a situation into account may arise through my lack of experience of situations of this kind. This is not to say that lack of experience always prevents a person from exercising good judgement, nor that the judgement of a more experienced person cannot be worse. The claim, rather, is this: if two people, neither of whom can be faulted in the first four respects, differ in their judgement about the moral character of an object, and the difference is to be explained by the lack of experience of one of the parties to the disagreement, then that is a good reason for preferring the judgement of the more experienced person.

This gives us a list of five characteristics of good moral judgement: why accept it? There are two good reasons. The first is our convergence on these factors when each of us identifies the sources of his own acknowledged failures of moral judgement. When a person accepts that his own earlier moral opinions have been improved, how does he explain the faultiness of his earlier judgement? The factors we converge on giving in such explanations at least include the five just identified. Secondly, the account of good *moral* judgement which they suggest corresponds to our account of good judgement concerning other, non-moral, subjects.¹⁰ To characterize a good judge of perceptual, aesthetic, economic, or other matters, we should give a list that follows the same pattern. In each case, a good judge must be someone who gives full attention to all the features relevant to the kind of assessment in question. But in no case does this suffice for good judgement. The relevant features must be attended to in the right way, and assigned the right relative importance. And failures to do this will correspond to the other failures we identified in the moral case. Inferential error or inconsistency will vitiate judgement in any of these domains. Inexperience will in each case count against my judgement, if it is the source of my disagreement with faultless judges. And corresponding to impartiality and sympathy, there is in each other domain a pair of analogous characteristics, one pertaining to the absence of bias, and the other to a distinctive form of interest or engagement appropriate to the subject matter. In the case of perceptual judgement, these are the normal functioning of one's sense-organs and normal external conditions; in the aesthetic case, they can plausibly be styled freedom from prejudice and "delicacy of imagination";¹¹ in the economic case (or, for that matter, in the predictive application of any science), we

¹⁰ As Mark Nelson reminds me, Rawls makes this point in supporting his similar list of characteristics of "competent moral judges": John Rawls, "Outline of a Decision Procedure for Ethics," *Philosophical Review* 60 (1951), pp.180-81.

¹¹ David Hume, "Of the Standard of Taste," in *Essays, Moral, Political, and Literary*, T.H. Green and T.H. Grose (eds.) (London: Longmans, 1907), pp.266-84, at p.272.

can cite a freedom from vested interest and a concern with understanding the relevant features of the case in the light of the best explanatory theory.

Why should continuity with judgement of other kinds count in favour of an account of good moral judgement? Because of the plausibility of characterizing moral judgements as objective, at least in the limited sense that when a person judges that a given object is M, she is committed to thinking two things: that it is M independently of her thinking so; and that anyone who thinks it is not M is mistaken. Just what kind of metaphysical view is required to vindicate moral objectivity is of course controversial. However, provided only that moral judgement does in fact have this objective character, the concern to avoid *distortions* of judgement will be central to our conception of good judgement. This concern is shared by the other domains in which we speak of good judgement, and explains why we should expect to find an appropriately corresponding structure to the account of good judgement in each domain.

V: THE REGRESS-PROBLEM

There are defensible claims to be made about the distinguishing characteristics of good moral judges. I have offered an initial list of five such features, but only to emphasize the general point that there are two plausible sources of support for an account of this sort. Now let us return to our problem. How can the virtue-ethicist use this to do any epistemological work? At first, the dilemma of Section III seems still to be firmly in place. On the first horn, the claim

All and only people who base their judgement on full, sympathetic and impartial attention to all and only the morally relevant features of a given object, who commit no inferential errors or inconsistencies, and who cannot be faulted for inexperience, are good judges of the moral character of that object

seems false, for precisely the reason foreshadowed in Section III: the list of five features does not give us a fully determinative rule for good moral judgement. Even if we specified which of the many different possible conceptions of sympathy and impartiality a

good judge must possess, that still would not guarantee a judge who assigned the right relative importance to the various morally relevant features of objects. Evidently, a good moral judge is not simply someone bearing the five characteristics; it is the *right kind* of person bearing them. These characteristics provide at most prima facie rules for someone's being a good moral judge. And the virtue-ethicist's earlier argument gives us reason to be sceptical about whether there is any further, fully determinative rule specifying the conditions under which someone is a good moral judge. To see when someone with this list of characteristics — or any further, improved list — counts as a good moral judge, I need — not an independent rule to apply, but — good moral judgement. And this impales us, it seems, back on the second horn of the dilemma: the regress-problem.

However, this problem can now be answered. To be warranted in holding a given moral opinion, I do not need a warrant for a complete set of premises entailing the conclusion that good moral judges would arrive at that opinion. All I need is an undefeated prima facie case for this conclusion. That will give me undefeated reasons for thinking what I do, and that is what warrant for judgement amounts to. A set of prima facie rules for good moral judges will supply me with such a case, provided it is reasonable for me to believe two things: first, that others' disagreement with me is explained by their being flawed in the respects mentioned in those rules, and secondly, that my own judgement is not so explained. (Without the second condition, the possibility would remain open that no-one's judgement is warranted.)

The regress-problem looked at first like an a priori demonstration of the impossibility of ever warranting a moral opinion on the virtue-ethicist's view. We have found that it is not this. What it does still suggest, however, is that a moral opinion of mine will be unwarranted unless I can satisfy the two conditions just given. If I cannot explain my opponents' disagreement with me in terms of the listed characteristics of good judgement, it still remains possible that their judgement is inferior to mine; but I cannot (according to the virtue-ethicist's view) be *warranted* in thinking that it is inferior, because any attempt to supply that warrant falls foul of the regress-problem. I could only warrant my claim

that I am a superior moral judge if I could warrant the claim that I am a superior judge of moral judges.

But given the extent and depth of actual moral disagreement, this may still seem to lead to moral scepticism. In particular, many moral disagreements about the moral value of an action themselves rest on disagreements about the moral relevance of features of the action. Agents whom I am inclined to fault for their lack of sympathy or impartiality often have a different conception of the kind of sympathy or impartiality that is morally appropriate. If my resources for faulting the judgement of people who disagree with me stop at the five simple features mentioned above, it seems they do not extend very far. And if faulting their judgement is required in order for my own to be warranted, then that seems to leave me with practically nothing that I *am* warranted in holding. There may not be an a priori argument for the impossibility of warranted moral judgement; but it may seem that it is hardly ever actually achieved.

VI: AVOIDING SCEPTICISM

The virtue-ethicist's view avoids this bleak conclusion. There are four general grounds for thinking this, the last being the most far-reaching.

First and most obviously, many of the most stubborn moral disagreements are based on background empirical or metaphysical disagreement. Here, the case just sketched for scepticism fails, since faulting my opponents does not require faulting their *moral* judgement. I may accept, for instance, that a religious fundamentalist with whom I disagree is a morally sensitive person, bearing all the characteristics of a good moral judge, without thereby having to withdraw my own opposing moral opinion. I should only do that if I come to doubt my warrant for holding the different religious beliefs I do.

The second point is that standards of warrant are not standards of proof. To be warranted in holding my opinion in the face of others' disagreement, I must be able reasonably to attribute their disagreement to a flaw in their moral judgement; but that does not mean being able to demonstrate that anyone who does not see this is confused. After

all, I can be warranted in believing something while others are warranted in taking the opposite view. Thus, if, in my experience, people who think and act like you tend to be, say, self-favouring, then that is a reason for me to doubt your judgement. This is evidence that you are flawed in respect of one of the features of good judges — evidence which falls well short of proof, but which may still make this the most reasonable thing for me to believe.

Thirdly, and less obviously, I can ask to what extent my opponents take account of my disagreement with them. Grounds for thinking they are failing to do this are themselves grounds for faulting their judgement. Why? On the virtue-ethical view, your opinion is warranted to the extent that you are warranted in faulting the judgement of those who disagree with you, and not faulting your own. Others' disagreement gives you a reason to doubt your own judgement, unless you can fault theirs. But if so, then someone who does not take into account others' disagreement is overlooking a morally relevant feature of the object being judged — namely, its being the object of a contrary opinion in someone whose judgement she cannot fault. And if she is overlooking a morally relevant feature, her judgement is flawed.

But even if I cannot fault my opponents' judgement on any of these grounds, it still does not follow that I am unwarranted in judging as I do. For a further simple point should be made — one with a clear analogue in the case of empirical judgement. If I cannot fault someone else's contrary judgement concerning a subject on which both of us cannot be correct, that gives me a reason to re-examine my own. But if, upon appropriate re-examination, I reasonably fail to find anything wrong with my own judgement, then surely I am warranted in thinking what I do. If it still seems to me that *p*, and there is no feature of my judgement that should lead me to doubt it, then I am warranted in believing that *p*. Does this mean abandoning the virtue-ethical test for warranted judgement, and allowing that I can be warranted even when I cannot fault the judgement of those who disagree? No. For the warrant I have for my own judgement itself gives me a case for faulting theirs. If an action strikes me as wrong, and upon appropriately re-examining my judgement it is reasonable for me to find no fault with it, then that is itself a reason —

clearly, a defeasible one, but in these circumstances the best I have — to think that it is my judgement that is right, and *therefore* that my opponents' is flawed.

For these four reasons, then, it is a mistake to think that the virtue-ethical view set out above is a recipe for moral scepticism. However, the position that has emerged needs, finally, to be defended against the contrary worry: that warrant is too easy to come by, not too difficult. Surely, it is often the case that neither party to a moral disagreement can initially fault the other in terms of our simple list of features of good judgement. If, whenever this happens, a second look at your own judgement suffices to warrant it, then it may seem that there is too little that will actually count as *unwarranted* moral opinion.

To respond to this, let us return to two questions raised earlier: the question of the basis for warranting judgements about morally relevant features, and the question of the relationship between “virtue ethics” as conceived here, and ethical theory.

VII: MORALLY RELEVANT FEATURES AND THEORY

Mentioning full attention to all and only morally relevant features of the object of judgement seemed indispensable to our list of characteristics of good moral judges. However, the obvious question it invites concerns how claims concerning moral relevance are themselves to be warranted. For moral disagreements can be disagreements about whether a feature is morally relevant. And, very commonly, they are disagreements about how much relevance a feature has, and its importance relative to other morally relevant features. How do I warrant the thought that I am the one making the correct judgements of moral relevance?

Now one way of answering this has just been given. If it seems to me that a given feature has a certain moral relevance, and if I reasonably fail to find fault with my judgement upon appropriately re-examining it in the light of others' disagreement, then that warrants me in holding the opinion I do, and in faulting my opponents' judgement. This, as we just saw, provokes a worry that it is too permissive, leaving us with too little *unwarranted* moral opinion.

A less permissive answer would be available, however, to someone who produced a plausible ethical theory. Such a theory would say which features are morally relevant, what *makes* them relevant, and the relations they bear to each other, as parts of a unified and systematic moral outlook. If unity and systematicity enhance the credibility of a person's judgements,¹² then my warrant for the moral opinions I hold will be strengthened to the extent that I can integrate them into a theory of this kind. And I would be entitled to claim a stronger warrant for my opinions than my opponents', if they lacked that theoretical support; so the permissive conclusion would be resisted.

What of the argument against ethical theory in Section II, though, accusing it of a faulty conception of the relation of rules to moral justification? The reply turns out to be simple. Most ethical theorists can readily disavow the aim of producing the kind of fully determinative rules targeted by that argument. They can deny that this is the status of the fundamental rules they identify ("Right actions are those maximizing well-being", "An act is wrong if its performance under the circumstances would be disallowed by any system of rules for the general regulation of behaviour which no-one could reasonably reject as a basis for informed, unforced, general agreement",¹³ and so on), and can embrace the idea that good judgement is needed for their correct application. This leaves plenty of room for relations of systematicity and unity to hold between the judgements

¹² One kind of anti-theoretical view denies this: it draws attention to the fact that features of moral relevance are not only reasons for judgement; they are also, and are primarily, reasons for *action*; and reasons for action (it holds) are not subject to the canons of systematicity and unity by which theory-construction is governed. I discuss this view in "Practical Theory," in Cullity and Gaut (eds), *Ethics and Practical Reason* (Oxford: Clarendon Press, 1997).

¹³ T.M. Scanlon, "Contractualism and Utilitarianism," in A. Sen and B. Williams (eds), *Utilitarianism and Beyond* (Cambridge: Cambridge University Press, 1982), p.110.

theorized in relation to these fundamental rules, yielding the epistemological conclusions just set out.¹⁴

It makes sense, then, to *look* for an ethical theory and, if I succeed, to claim a warrant for my opinions that my opponents' do not possess. But there is no guarantee that there is any such theory to be found. For while the virtue-ethicist's argument fails to undermine the search for an ethical theory, it *does* still show that such a theory is not required for warranted moral judgement. Section II did undermine the thought that we need a theory codifying independent rules which are necessary for moral justification. And Sections III-VI have shown that a satisfactory account of the warrant moral judgements can possess will be available even if the search for a theory fails. Without a moral theory, we do seem to have a more permissive view about warrant.¹⁵ But it is not clear what should be thought objectionable about the permissive view. All it deprives me of is the ability to say that those who disagree with me are unwarranted in doing so. It does not deprive me of a warrant for thinking that those who disagree with me are incorrect; and it is only this that I require if I am to warrant my own opinions.

As developed here, then, the virtue-ethicist's argument leads us to the conclusion that **an** anti-theoretical but non-sceptical position remains a possibility. What must be abandoned, though, is the original suggestion that this is the *only* possibility: perhaps the search for an ethical theory will succeed. What is at stake, I have argued, is the permissiveness of the resulting moral epistemology. My success in finding a theory would deprive me of a warrant for judgements that would otherwise be warranted.

¹⁴ For defences of ethical theory in this spirit, compare T.M. Scanlon, "The Aims and Authority of Moral Theory," *Oxford Journal of Legal Studies* 12 (1992), Sect.III; Samuel Scheffler, *Human Morality* (New York: Oxford University Press, 1992), Ch.3.

¹⁵ Just *how* permissive a view is generated is a further large question that will have to remain unsettled here. It depends on how often it really will be reasonable to find my own moral judgement to be faultless.

The first of the conclusions announced in Section I has been drawn: the virtue-ethical attack on ethical theory fails. And the second should also now be evident. There is no sense in which the theorizing consistent with virtue-ethics will supply us with a distinctive “virtue theory”: the theories I have been discussing are simply the familiar candidates. Virtue-ethics neither undermines nor produces ethical theory. What it has given us, however, is a strong framework for moral epistemology, which makes an anti-theoretical but non-sceptical position a clear possibility — the one we are left with if our attempts at theory-construction fail.