

WILEY

AS

Particularism and Moral Theory

Author(s): Garrett Cullity and Richard Holton

Source: *Proceedings of the Aristotelian Society, Supplementary Volumes*, Vol. 76 (2002), pp. 169-209

Published by: [Wiley](#) on behalf of [The Aristotelian Society](#)

Stable URL: <http://www.jstor.org/stable/4106967>

Accessed: 03/07/2013 01:49

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Wiley and *The Aristotelian Society* are collaborating with JSTOR to digitize, preserve and extend access to *Proceedings of the Aristotelian Society, Supplementary Volumes*.

<http://www.jstor.org>

PARTICULARISM AND MORAL THEORY

by Garrett Cullity and Richard Holton

I—Garrett Cullity

PARTICULARISM AND PRESUMPTIVE REASONS

ABSTRACT Weak particularism about reasons is the view that the normative valency of some descriptive considerations varies, while others have an invariant normative valency. A defence of this view needs to respond to arguments that a consideration cannot count in favour of any action unless it counts in favour of every action. But it cannot resort to a global holism about reasons, if it claims that there are some examples of invariant valency. This paper argues for weak particularism, and presents a framework for understanding the relationships between practical reasons. A central part of this framework is the idea that there is an important kind of reason—a ‘presumptive reason’—which need not be conclusive, but which is neither *pro tanto* nor *prima facie*.

I

T*wo Kinds of Particularism.* Moral particularists and their critics debate two issues. The two are related, but not as closely as is often thought.

One issue concerns the existence and nature of moral principles. Are there any correct general principles that can serve to justify moral judgements? The kinds of moral judgements we are ultimately interested in reaching are overall verdicts about the objects of moral assessment—judgements about whether actions are right or wrong, whether a person or a way of living is virtuous or vicious, whether a state of affairs is (all things considered) good or bad, and so on. Call these ‘verdictive judgements’. The question is then whether there are any general principles non-trivially linking verdictive moral properties—properties such as the rightness of an action or the viciousness of a person—to other properties. Or to put the question in a more metaphysically cautious way, we can ask whether there are any general principles linking the application-conditions of verdictive moral terms to the application-conditions of other terms.

The second issue concerns whether, if a consideration counts as a good reason for an action in one situation, it must do so in others.¹ To take the particularists' leading example, it seems that normally the fact that I would enjoy something is a good reason for me to do it, and the fact that you would enjoy doing something is a good reason for me to help you to do it. However, they maintain, this is not always true. The enjoyment of cruelty or humiliation does not count in favour of an action at all: on the contrary, it actually counts against it. The normative 'valency' of enjoyment changes from one context to another: in many contexts, it counts as a reason for action, but in some it does not.

Self-styled particularists typically make a claim about each of these issues.

- (i) There are no exceptionless, finite general principles specifying the descriptive conditions under which a moral verdict is justified.²
- (ii) At least some considerations have variable normative valency.

However, these two claims are logically independent, and it is easy to find adherents of one of them who deny the other. Roger Crisp (2000) and Joseph Raz (2000) endorse (i) while arguing against (ii). And Frank Jackson, Philip Pettit and Michael Smith (2000) reject (i) while being prepared to accept (ii).

It makes sense, therefore, to distinguish particularism about principles from particularism about reasons.³ Each comes in different strengths. Weak particularism about principles⁴ is the

1. A related but different issue is whether, if a consideration counts as a reason for a given *evaluative verdict* in one situation, it must do so in others. Jackson, Pettit and Smith (2000), pp. 96–9 and Little (2000), p. 280 discuss this further issue. Raz (2000), pp. 58–61 and Crisp (2000), pp. 32–42 discuss the one in the text. Dancy (1993), Ch. 4 and McNaughton and Rawling (2000) move between the two.

2. A proponent of (i) owes us an account of the distinction between the descriptive and the evaluative. It seems to me that a satisfactory account can be given by saying that evaluative terms are those the assertoric use of which expresses a judgement of goodness or badness, and then understanding descriptive terms by contrast. The distinction between descriptive and evaluative *properties* can then be understood derivatively.

3. Compare Crisp (2000), who also distinguishes these two views from particularism about motivation.

4. Examples of weak particularists about principles are McNaughton and Rawling (2000), Raz (2000), Crisp (2000) and Nussbaum (2000).

acceptance of (i)—the view that there is no finite list of conditions D yielding any correct general principle of the form:

Under descriptive conditions D, the correct moral verdict is V.

Someone who thinks this, however, might well think that there are principles of a less ambitious kind—principles of the form:

Descriptive conditions D always count in favour of moral verdict V.

Such principles propose that there are general, descriptively specified *pro tanto* reasons—reasons that always count in favour of a given moral verdict, although perhaps not always decisively so. A stronger kind of particularism about principles denies the existence of correct principles of this more modest form.⁵ And the strongest kind of particularist about principles denies that there are any true, non-trivial general principles of the form:

The presence of R supports moral verdict V.⁶

On this view, even ‘thick’ moral properties such as cruelty have variable valency: it is not simply that cruel actions are sometimes right because there are stronger reasons in their favour; for some cruel actions, their cruelty does not count against them at all.

Turning to particularism about reasons, we find that this equally comes in different strengths. Weak particularism about reasons is claim (ii): the claim that there are some considerations that are reasons in some contexts but not others.⁷ A stronger view is that this is true of every descriptive consideration: every descriptive consideration can be a reason, but all descriptive reasons sometimes change their normative valency.⁸ And once more, there is a stronger view still: the view that all non-verdictive reasons (even those supplied by ‘thick moral properties such as cruelty) sometimes change their normative valency.⁹

5. E.g. McNaughton (1988), Ch. 13, and Little (2000).

6. E.g. Dancy (1993), Chs 4–6.

7. I prefer to talk of reasons as ‘considerations’ rather than ‘facts’ because I think of them as linguistic entities, and not the states of the world these entities represent. My reason for this is that the content of a state of the world is always fully specific; but the content of a reason (as I shall argue in Section III) is not.

8. E.g. Little (2000).

9. See again Dancy (1993), Ch. 4.

These two issues are connected. One connection between them is this. Some reasons to ϕ are reasons for holding that ϕ ing is morally right: call these 'moral reasons'. Then for any descriptive consideration D which is a moral reason of invariant valency, there will be an exceptionless principle of the form:

D always counts in favour of the moral rightness of action.

Strong particularism about moral reasons entails and is entailed by particularism about principles.

To support their view, strong particularists about both reasons and principles commonly appeal to an argument which, in condensed form, goes as follows.¹⁰ The evaluative is 'shapeless' with respect to the descriptive. That is, no finite disjunction of our descriptive concepts has the same extension as any evaluative concept.¹¹ Wittgenstein taught us to reject the prejudice about rationality of thinking that a practice of concept-application must have an independently articulable backing rule in order to count as rationally constrained. In order for me to count as genuinely going on in the same way in applying a concept, it need not be the case that there is any way, independent of the use of that concept, of spelling out the conditions under which I count as successfully doing so. If so, we should deny that there is any reason to expect that evaluative verdicts should be backed by principles linking them to the descriptively specified conditions in which they obtain. Moreover, once we see that the evaluative and descriptive are independent in this way, we should press a further question. Why think that the contribution that any descriptive characteristic makes to the evaluative character of the situations in which it is present is independent of the rest of the context in which it is found? In the absence of any compelling reason for thinking this, we won't have *pro tanto* principles linking the descriptive and the evaluative either. That leaves us

10. See McDowell (1979) and (1981); Dancy (1993), esp. Ch. 5, Section 4; and Little (2000), esp. Section II; also, for discussion, McNaughton and Rawling (2000), Section II.

11. It is hard to see how they could have a good argument for thinking there *could not* be descriptive concepts with the same extension as any evaluative concept; but as I read it, their argument relies only on the fact that we do not in fact have such concepts. I am grateful to Frank Jackson for a very helpful correspondence on this point.

with a view on which reasons emerge holistically, from the entire conjunction of descriptive features that make up the evaluative character of a situation: the reason-giving character of a descriptive feature does not attach to it in an atomistic, context-independent way.

In this paper, I defend a version of weak particularism about reasons. I shall argue that the normative valency of some descriptive considerations varies, but others have an invariant normative valency. As far as I am aware, this view has not been defended by other contributors to the debate. This is surprising in one way, because it is such a natural thing to say about the examples I shall consider. In another way, however, the rarity of this view is not so surprising. For it requires arguing against opponents from two different directions, and that may seem to make it an unstable position. I need to respond to arguments that a consideration cannot count in favour of any action unless it counts in favour of every action. But I need to do so without resorting to a global holism about reasons, if I am to claim that there are some examples of invariant valency. Arguing against these two opposing views will occupy Sections III and IV, respectively. What will emerge from this, in Section V, is a framework for understanding the relationships between practical reasons. A central part of this framework is the idea that there is an important kind of reason which need not be conclusive, but which is neither *pro tanto* nor *prima facie*: I shall call it a 'presumptive reason'.

II

The Content of Reasons. Particularists and anti-particularists about reasons can agree about the conditions under which a person has a reason. They can agree, for example, that when an innocuous action of mine would be enjoyable, there is a reason for me to do it, and that the fact that an action would produce sadistic enjoyment does not amount to a good reason for doing it. Their disagreement concerns the *content* of the reasons that we have. According to the particularist, the content of the reason in the first, innocuous case is simply *that I would enjoy doing this*. In my innocuous circumstances, this consideration is a good reason for acting; but in the sadist's different circumstances, the same consideration is not a good reason. The anti-particularist's

view, by contrast, is that the description of the reason in the first case is too simplistic. The unqualified fact that I would enjoy doing something is not enough to give me a reason: after all, if it is the wrong kind of enjoyment then I do not have a reason. In the situations where I do have a reason, the reason must be at least *that I would non-sadistically enjoy doing this*. Admittedly, it would be pragmatically odd to offer this as a reason for going to the pub. But speaking with strict accuracy is often pragmatically odd. A *complete* statement of the reason that does count in favour of my action will have to include the qualification.

Of course, qualifying the content of the reason by simply excluding sadistic motivation looks inadequate. If you do not think that there is a reason to pursue every form of enjoyment, you are unlikely to think that sadistic enjoyment is the only exception. There are forms of malice that it seems incorrect to describe as sadism; forms of enjoyment of the spectacle of harm to others that involve no malice, enjoyment of destruction which involves harm to no one, and so on. It looks as though the only succinct way to capture these exceptions will be to use an evaluative term, and say that when there is a reason, it is *that I would innocuously, or morally permissibly enjoy doing this*.

One strong-looking argument for the anti-particularist view of the content of reasons can be put like this. Surely a fact can only be a reason for a given action if the obtaining of that fact is sufficient to make it the case that there is a reason to perform that action. However, the fact that an action would be enjoyable cannot be sufficient to make it the case that I have a reason to do it, if there are kinds of enjoyment that are not reason-giving. Therefore, when there is a reason to do something enjoyable, the reason cannot simply be *that it is enjoyable*. In order to mention the whole of the reason for doing it, we need to mention the kind of enjoyment that it involves, specifying that it is one of the kinds that is reason-giving, rather than one that is not.

For a second argument to the same conclusion, we can turn to Joseph Raz.¹² Raz's argument against particularism about reasons begins with what he calls 'the intelligibility of value'—the idea that

12. Raz (2000), esp. Section 3.

there is nothing 'arbitrary' in the domain of value.... There is an explanation for everything, an explanation for why what is good is good, what is bad is bad, etc.¹³

Given this, it is not enough simply to say that the same consideration might be a reason in one situation but not in another: there must be some difference between the two situations which *explains* the evaluative difference. But that difference, according to particularists, is not itself part of the reason: if it were, that would mean there were different reasons in the two situations. The upshot is that, on their view, 'not everything relevant for the evaluation of an action is part of the reasons for or against the action.'¹⁴ But that looks wrong, for the notion of a reason for an action just is the notion of what counts in its favour.

Roger Crisp sets out a third argument for anti-particularism about reasons.¹⁵ He points out that it is not difficult, in the sorts of cases discussed by particularists, to give an explanation of the differences between the situations in which they claim that the normative valency of a consideration varies. In each case, a transparent explanation is given by using virtue-terms. Sometimes, doing what is enjoyable is prudent; at other times, it is cruel. Sometimes, lying is dishonest; at other times, it isn't. Sometimes, doing what is illegal is unjust; at other times, it is required by justice. But that is to say that the *ultimate* reasons for action in each of these pairs of cases are given by these evaluative considerations, the normative valency of which does not vary. It may make pragmatic sense to cite as a reason for the wrongness of an action the fact that it was a lie, but when lying is wrong an explanation is readily available of why it is wrong, and that explanation amounts to the provision of the ultimate reason why it is wrong.

III

Reasons, Motivation and Explanation. I think that these arguments fail. To show this, I shall give a three-stage argument in support of particularism about reasons; then I shall explain what

13. Raz (2000), p. 50.

14. Raz (2000), p. 60.

15. Crisp (2000), pp. 32–42.

was wrong with each of the three anti-particularist arguments just given. The first stage of my argument connects normative reasons to motivation. The second connects motivation to a particular kind of explanation: *self*-explanation. The third then connects explanations with expectations of normality.

There are various different claims concerning the conceptual connection between normative practical reasons and the motivation of rational agents that have been called 'internalism about reasons'.¹⁶ One claim which seems correct is this: a reason for me to ϕ in circumstances C must be a consideration my rational orientation towards which could motivate me to ϕ in C, were I acquainted with the facts about C.¹⁷ That is, my being aware of this reason and rational guidance by that awareness could explain my ϕ -ing. Given the nature of the concept of a reason, and the concept of rationality, it is hard to see how this connection could sensibly be denied. A normative practical reason is a consideration that counts in favour of my performing the action for which it is a reason: it is a consideration which makes it the case that I *ought* to do it—or, to use Allan Gibbard's phrase, a consideration which shows why my doing it *makes sense*.¹⁸ And rationality, although it need not involve being successfully guided by the reasons for acting that apply to you (for you might be rationally mistaken about these), at least involves being appropriately guided by what you are *warranted* in treating as the reasons for acting that apply to you. If that is right, then the only way in which there could be a reason for me to ϕ that could not explain my ϕ -ing when rationally oriented towards it would be if I could not be warranted in regarding it as a reason. But how could a consideration that I could never be warranted in regarding as a reason really count in favour of my performing any action? It could never *make sense* for me to act on such a

16. For further discussion of a claim closely related to the one I concentrate on here, see Parfit (1997). For some of the others, see the taxonomy in Audi (1997); also Darwall (1983), esp. Ch. 5.

17. One kind of 'externalism' about reasons is McDowell's (1995) view that there can be reasons for me to act which are such that, given my actual motivations, there is no rational process through which I could come to be motivated to act on them. This would have to occur by a process of non-rational 'conversion'. Notice that even this view is consistent with the 'internalism' proposed in the text. McDowell's reasons are still reasons I *could* (if 'converted') be rationally motivated by.

18. Gibbard (1990), Ch. 1.

consideration, no matter how rational I was. But a 'reason' that we could never be warranted in treating as a reason looks like no reason at all. It would not be worth discussing, since it would not be something that it made sense for a rational person to care about. And if not, the claimed connection between reasons (the only reasons it makes sense for us to be concerned with) and rational motivation follows: it is a necessary condition of a consideration's being a normative reason for me to ϕ that my rational orientation towards that consideration could motivate me to ϕ .¹⁹

This is a conceptual, not an empirical claim. It leaves it open that there are no considerations that meet this condition, and therefore no reasons. And it leaves it open that there are reasons, but we are never rationally oriented towards them. What it claims is this: for any consideration to satisfy the concept of a reason, and for any person to satisfy the concept of a rational agent, they must meet this condition. Notice also that it does not require the principle that 'ought' implies 'can'. For it leaves it open that even when I *cannot* be rationally motivated, I ought to do the things I would do if I were rationally motivated.

Now for the second stage of the argument. When we say what motivated an action, we are offering a certain kind of explanation of it: an explanation in terms of the agent's aims in performing it. (I might explain what you are doing by saying that you're confused; but it would not ordinarily make sense to say that you were *motivated* by confusion.) Normally, specifying those aims will involve saying what the agent thought of as counting in favour of the action.²⁰ That is, it will involve saying what the agent regarded as the normative reasons that made the action worthwhile. Motivational explanation, then, is normally explanation in terms of the agent's judgements about normative reasons. Suppose we ask this, though: explanation by whom, to whom? The answer is: in the first instance, explanation by me,

19. 'My rational orientation towards that consideration could motivate me to act' is more awkward than, 'I would be motivated to act on that consideration if I were fully rational'. It deals better, however, with situations where I have reasons to correct my own irrationality. I would not be motivated to act on those reasons if I were fully rational. But acting on those reasons is what I will do insofar as I am rational in responding to my own irrationality.

20. I say 'at least normally' in order to allow for the possibility of the kind of counter-rationally motivated action discussed in Stocker (1979) and Velleman (1992).

the agent, to myself. If my action is motivated by the pursuit of some goal then that is to say that, at some level (the level at which I am motivated), I am seeing my action as the pursuit of that goal. At that level, I am giving an account to myself of why I am doing *this*, rather than something else. That account corresponds to the content of my goals in acting. When a third person explains my action by saying what motivates it, he is therefore giving an account of how I am explaining the action to myself.

When I say this, I am not saying that we have to privilege an agent's own reports of his motives. We can be dishonest or even self-deluded about our motives. My claim is simply that in any case of motivated action, the correct account of what motivates me takes the form of attributing to me the pursuit of a goal through the action, and that my pursuing a goal through the action amounts—perhaps subconsciously—to my giving an explanation to myself of why I am doing what I do. This may not stop me from giving other, contradictory explanations of the same action to myself and others, and from wrongly believing that it is the more palatable goals that are really motivating my action. Suppose I think I'm being kind to you, but I'm really trying to dominate you by making you indebted to me. Saying this only makes sense if the action is explained by my seeing it, at some level, as enabling me to dominate you. The claim is that, at that level, I am explaining the action to myself as an act of domination.

The third stage of the argument concerns the nature of rational explanation. Clearly, the explanation it is rational to give of any phenomenon is relative to background expectations of normality. If the trees next to the vineyard normally flower before the vintage, it may be rational to appeal to the fact that they did not in order to explain why the vineyard suffered a lot of bird damage this year; but if not, it will not. If the trees never flower before the vintage, it may be true that had they done so, there would not have been a lot of bird damage; but it does not follow that their not flowering before the vintage should be included in an explanation of the bird damage. In general, it is fallacious to reason that if A would not have happened in the presence of B, the absence of B should figure in a rational explanation of why A happened.

Putting these three stages together, we have an argument for the particularist claim about the content of reasons. A reason for me to ϕ must be a consideration my awareness of which could give a motivational explanation of my ϕ -ing, if I were rationally oriented towards that consideration and aware of the other facts. But a motivational explanation is primarily a self-explanation; so a reason for me to ϕ must be something I could invoke in explaining my ϕ -ing to myself, insofar as I am rational.²¹ But the content of a rational explanation is relative to background expectations of normality. If so, the background expectations of normality it is rational to have will constrain the content of normative reasons, in the following way. The absence of those considerations that would defeat the presence of a reason will not normally be part of the content of my reason. For, in straightforward circumstances, they are not part of what I can rationally invoke in explaining my action to myself. When I have an opportunity to do something enjoyable, in a situation in which questions of permissibility do not normally arise, it will simply be the fact that the action would be enjoyable that features in a rational explanation to myself of what I am doing. It might be true that were my enjoyment morally objectionable, I would not have a reason to do what I am doing. But that does not mean that its being morally unobjectionable is part of the reason I do have. For its being morally unobjectionable will not normally be part of a rational explanation to myself of what I am doing; hence not part of the goal that motivates me; and if it could not be part of what motivates me insofar as I am rational, then it cannot be part of the reason there is for me to act.²²

21. This is a necessary condition on something's being a reason; not a sufficient one. It might be rational for me to explain my action by appealing to something that is not a reason at all (as Frank Jackson pointed out in discussion).

22. My understanding of what is normal will change over time. Does that mean the content of the reasons there are for me to act will change? Suppose I begin by being sensitive to whether my enjoyment of your company is morally objectionable (I've been told you're an unsavoury character). Then, as I get to know you, the fact that this enjoyment is unobjectionable becomes part of my background expectation. Does that mean that, although the unobjectionableness of the enjoyment was part of the reason there was for me to spend time with you initially, that ceases to be true?

No. My claim is that a reason must be something I could be rationally motivated to act on, *given knowledge of the facts*. If the background facts themselves change, my reasons will change: that does not seem an embarrassing result. But this does not imply that the reasons I have will change simply because my rational expectations change. Thanks to Philip Pettit for pressing me on this.

The previous section gave three arguments for the opposite conclusion. We can now say what was wrong with them. The first maintained that the fact *that I would enjoy doing this* is not sufficient for the existence of a reason for action. There is a clear sense in which this is true: this fact is not sufficient for the existence of a reason in *all* circumstances. However, given the argument I have just presented, we should still say that this fact is sufficient for the existence of a reason for action in most circumstances, where the background is uncomplicated. We must be careful not to commit a counterfactual fallacy about practical reasons that corresponds to the fallacy about explanation identified earlier. When there is a reason for me to ϕ , but there would not have been a reason for me to ϕ had consideration C been present, it does not follow that the absence of C must be part of the reason for me to ϕ .

Raz's argument involves a dubious inference of a different kind. He begins by observing, correctly enough, that on the particularist view, if a consideration counts as a reason in one situation but not another, there must be some further feature of the situation which is not part of the content of the reason, and which explains the difference. But he moves from there to supposing that this explanation of the evaluative difference cannot be part of the reasons for and against the action for the particularist, and complaining that this drives an unsatisfactory wedge between reasons and what explains evaluative nature. This does not follow, though. The evaluative difference, on the particularist's view, lies outside the content of the reason whose variable valency is being explained. But that does not mean that it lies outside the content of any practical reason. Indeed, in the case we have been considering, we can see that this is not so. Enjoyment may fail to be a reason for an action when it is malicious. On the particularist's view, we should say that non-maliciousness is not part of the reason for doing enjoyable non-malicious things. But it still makes sense to say that their being malicious *is* part of the reason for *not* doing things that are malicious. Facts about whether something is malicious or not are not always part of the reason for doing enjoyable things, on the particularist view. But that does not mean that it has to say that the factors that explain the difference between the situations in which you ought to do what is enjoyable and the situations in which you

ought not are not part of the content of practical reasons at all.²³ Crisp's anti-particularist argument was that we can make sense of cases of apparently variable normative valency by appealing to invariant ultimate reasons, expressible using virtue-terminology. But our three-stage argument shows what is wrong with this. Crisp's 'ultimate reasons' are not the reasons that motivate rational agents: my reason for going to the pub is that I would enjoy it, not that it would be prudent. Maybe it *would* be prudent, and maybe (prudent person that I am) I wouldn't go if it were imprudent. But in a situation where issues concerning its possible imprudence do not arise, my going is not explained by my seeing it as prudent. Reasons are the considerations that motivate us insofar as we are rational; so it is the simpler considerations that supply our reasons in straightforward circumstances.

IV

Resisting Holism. I have argued for the first half of my weak particularism about reasons: the claim that there are some descriptive considerations that have variable normative valency. The other half is the claim that some have invariant valency.

This second claim is one that other particularists about reasons—writers like Dancy, Little, McNaughton and Rawling—all reject.²⁴ Their thought is that once we notice the variable normative valency of some descriptive considerations, and once we notice the shapelessness of the evaluative with respect to the descriptive, we will see that we need to embrace the general thesis that the way descriptive considerations function as reasons

23. Raz discusses at some length a different response that particularists might make to his argument: a reason for an action must be capable of being the reason why someone acted; but no one can be guided by all the evaluatively relevant factors that are present (p. 61). He replies that if reasons are objective, one can refer to them without understanding them fully (pp. 61–9). This does not seem decisive: we should certainly agree that I can *refer* to reasons without understanding them ('the reasons set out in the book I ought to read'); but what is less clear is how I can be *guided* by them without understanding them. However, it seems unwise for a particularist to take the line Raz is considering. For although particularists about principles do want to say that there is no finite descriptive account of the evaluatively relevant factors, they allow that there are evaluative terms capturing all the relevant factors, by which we can be guided.

24. Dancy (1993), Little (2000), McNaughton and Rawling (2000). In the vocabulary preferred by the latter, the position defended here is a version of 'fat intuitionism'.

is holistic. Descriptive considerations, on this view, are of the wrong kind to have invariant normative valency: the evaluative significance of any descriptive consideration depends on the rest of the context in which it is found.

This is a surprising view, since holism's leading claim about the normative status of descriptive reasons—the claim that *every* descriptive consideration has variable normative valency—is obviously wrong. Indeed, our discussion of why enjoyment should be thought to have variable normative valency supplies us with the most obvious example of a descriptive consideration with invariant normative valency: inflicting suffering on others for your own enjoyment. This always counts against an action. Indeed, a stronger claim looks plausible: such actions are always wrong. Even if someone deserves to suffer, and even if there are further reasons that make it very important to do what makes him suffer—even if he is an evil megalomaniac who needs to be harmed in order to save the world—it is wrong to make him suffer for your enjoyment. Maybe you should make him suffer because he deserves it, or because it will save the world; but it would always be wrong to make him suffer for your own enjoyment.

Finding this descriptive property was easy. Having noticed that enjoyment sometimes counts as a reason and sometimes does not, we have simply taken one class of cases in which it does not, and noticed that when this class of cases is described in a way that includes the agent's motivation, that generates a description of invariant normative valency. This will give us a recipe for generating other descriptive reasons of invariant valency: stealing for enjoyment, annoying someone for enjoyment, forcing someone to do something so that you can enjoy feeling superior to him, and so on.

Why should it be the case that some descriptive reasons have variable moral valency while others are invariant? The core of the explanation is this. Sometimes, in stating a descriptive reason, we are referring to a state, such as enjoyment, that can have different contents. I have argued that the particularists are right that often, the correct description of the reason for doing something enjoyable takes the simple, 'content-neutral' form: *I would enjoy this*. However, one of the ways in which the contents of a state such as enjoyment can differ is that they can have different

values. Many of the contents of enjoyment are not bad, but some are bad. And there are some ways of describing the content of a state such as enjoyment which make it *always* bad. If so, we can use the descriptions of those contents in order to generate further, more specific considerations—such as the consideration that this would be an action of harming someone for enjoyment—which amount to reasons of invariant valency. This is consistent with thinking that, in circumstances in which there is no reason for me to be raising the question of whether the content is bad, the reason should be stated in the simple content-neutral form.

This is suggestive; but we need to take it further. What other ‘content-neutral’ reasons are there? How do they relate to each other? And why should it be the case that some contents of a state such as enjoyment make it invariably bad? What I offer next is no more than a start towards answering these questions. But I think enough can be said to indicate the lines along which these questions can be answered.

V

Presumptive Reasons, Practical Norms, and Undermining. In Section III, I opposed Raz’s argument for anti-particularism about reasons. However, I think he is right about what he calls ‘the intelligibility of value’. It cannot simply be a brute evaluative fact that a certain consideration is a reason in one place and not in another: there must be a justifying explanation of why the two cases are evaluatively different. The idea of that which is supported by reason is to be fundamentally contrasted with the idea of that which is arbitrary.

Now we have at least the beginnings of an answer to the question why a consideration that gives a reason in one situation does not give a reason in others. ‘Content-neutral’ descriptive considerations refer to states that can have good and bad contents; and the difference in the values of these contents explains the difference in the normative status of those considerations in different situations. However, we cannot stop there. A concern with the intelligibility of value will resurface as the question: what explains why some contents are good and others bad?

Having opened this question, we might wonder what would count as a *complete* answer. What form must a justification of

evaluative claims take, if it is not ultimately to appeal to brute evaluative facts?²⁵ I do not propose to offer an answer to that larger question here. But it is possible at least to make a start on explaining why some contents of ‘content-neutral’ descriptive considerations are bad, and thus why they fail to provide reasons when they have those contents.

In order to do this, we should start by asking which other content-neutral descriptive considerations are reasons. When a consideration provides a normative reason for an action, I shall refer to the relationship of support that exists between the consideration and the action as a practical norm. And I shall use the following arrow notation to express the existence of a practical norm:

- (1) ϕ -ing would be enjoyable
 \downarrow
 ϕ -ing.

This is the norm we have been concentrating on so far. Other descriptive considerations that can favour an action give us other simple norms, for example:

- (2) ϕ -ing would be interesting
 \downarrow
 ϕ -ing.

(1) and (2) are naturally thought of as non-moral norms: someone who fails to follow them, the thought goes, is failing to be properly responsive to reasons, but they are not open to *moral* criticism.²⁶ However, we can extend our list of descriptive reasons to cite those that seem to provide us with moral norms:

- (3) ϕ -ing would get her X
 she needs X
 \downarrow
 ϕ -ing.

25. According to one line of thought, taking this problem seriously draws us towards a Kantian position. If answers to the question ‘Why?’ are to avoid either resting on an unvindicated claim about brute evaluative facts or leading to an infinite, uncompletable regress, this requires us to give an account of the formal nature of reasons, and derive from this formal account substantive conclusions concerning the particular reasons we have. See O’Neill (1992) and Korsgaard (1996).

26. Can moral and non-moral reasons be clearly distinguished, and if so, is the distinction important? I am not relying on any answer to these questions here.

- (4) she wants to know whether or not P
P is the truth
↓
asserting P.
- (5) others are relying on me to ϕ
↓
 ϕ -ing.
- (6) ϕ -ing would be forcing him against his will
↓
not ϕ -ing.
- (7) ϕ -ing would be lying
↓
not ϕ -ing.
- (8) I have promised to ϕ
↓
not ϕ -ing.
- (9) ϕ -ing would be illegal
↓
not ϕ -ing.

In (1)–(9), we have a plausible (if not exhaustive) list of descriptive reasons that are ‘content-neutral’. In each case, the argument of Section III supports the view that, in straightforward circumstances, good reasons for acting can have the simple, unqualified contents just listed. However, in each case, the descriptive consideration refers to a state whose content can be good or bad. Often, this is because it refers directly or indirectly to attitudes of a person that can be well- or badly-directed. We have seen this in the case of (1). Normally, enjoyment is good; but when it is enjoyment of what ought not to be enjoyed, it is bad. The same general point applies to (2)–(6). If I am only interested in something as part of a bad enterprise—I am interested in finding out more about how to harm or swindle other people, say—then the fact that I find it interesting is not a good reason: it is the wrong kind of interest to provide a good reason. The same goes for needing something, wanting to know something, or relying on me for something: these could all be serving a bad enterprise. And likewise, if someone’s will is directed towards

what is bad, forcing him to act against his will may cease to be something there is a reason to avoid.

Norm (7), concerning lying, gives us a subtler case of content-neutrality. Here, unlike (2)–(6), there is no direct reference to attitudes of another person that may be well or badly directed. However, the explanation of the variable normative valency of lying is similar. There is no reason not to lie to a murderer at the door because the evil nature of his enterprise means there is no reason not to deceive him. We have a reason not to deceive other people out of respect for their pursuit of their own ends. But sometimes, others' pursuit of their ends is not respect-worthy: indeed, sometimes we ought to frustrate it. Deceiving people in relation to *these* ends is not something there is a reason to avoid.

The reasons spelt out in (8) and (9) are content-neutral in a different way. Here, the normative variation in content is not a variation in the contents of people's attitudes or enterprises, but more directly a variation in the content of promises or laws. Promises and laws that are directed towards an evil purpose are ones that we have no reason to keep.

What we have started to compile is a list of norms associated with different virtues. In straightforward circumstances, (3) gives us the reason the recognition of which is characteristic of beneficence or kindness. (4) and (7) are the simplest manifestations of one sort of honesty: honesty-as-veracity. Another sort, honesty-as-fidelity, is found in (8). A reliable person is someone who often follows the norm set out in (5), a respectful person (6) and a law-abiding person (9). We might want to say that the *joie-de-vivre* associated with (1), and the curiosity associated with (2) are not *moral* virtues. But if we do, we should still be prepared to say that a good person—someone who is properly responsive to the reasons there are—is someone who is guided by (1) and (2), along with the other norms we have listed.

I do not want to suggest that every norm associated with a virtue contains a content-neutral descriptive reason, as (1)–(9) do. Very often, the content of the reason recognized in virtuous agency will be evaluative, even in the simplest cases. Thus a further kind of honesty involves following the norm:

- (10) ϕ -ing would be stealing
 \downarrow
 not ϕ -ing.

where stealing will have to be understood as an evaluative notion, since stealing is taking from someone not simply what they possess, but what they are entitled to possess. And norms for justice, conscientiousness and public-spiritedness include these:

- (11) ϕ -ing would result in distribution D of these goods
D is the fair distribution
↓
 ϕ -ing.
- (12) ϕ -ing would get her X
she has a right to X
↓
 ϕ -ing.
- (13) I have a duty to ϕ
↓
 ϕ -ing.
- (14) we all ought to be ψ -ing
we can only ψ if people like me ϕ
↓
 ϕ -ing.

These evaluative considerations can have different contents, but they do not count as 'content-neutral' in the sense employed here, since it is not natural to think of their contents as sometimes being bad.

I have said that these norms are 'associated' with virtues. But what, more precisely, is the association? Clearly, it is not that any virtue simply consists in following one of these norms. In the case of (10)–(14), we might want to say that possessing the virtues associated with them is at least in part a matter of recognizing the evaluative considerations they cite as always giving *pro tanto* reasons. However, it is not true that the considerations mentioned in (3)–(9) are *pro tanto* reasons wherever they obtain. So why think that these norms tell us anything important about the virtues with which they are (sometimes) associated?

I think that in each case there is a deep association between the norm I have indicated and the relevant virtue: each of these norms gives us the core of a virtue. To explain this, we need a

term to describe the kind of reason that is provided by a consideration such as that doing something would be enjoyable, or would be the fulfilment of a promise. These considerations, I have maintained, do not always give us *pro tanto* reasons. They sometimes do, but at other times, their status as reasons is undermined by further considerations (ones that show that my enjoyment, or my promise, is bad). But it would be too weak to say that they give us merely *prima facie* reasons: considerations that on first inspection appear to be reasons. They *are* reasons, unless undermined. We can mark this special status by giving them a label of their own. These considerations are *presumptive reasons*—that is, they are *pro tanto* reasons unless undermined. Being properly responsive to reasons requires recognizing this. And being morally good requires recognizing the considerations set out in (3)–(14), amongst others, as presumptive reasons.

This allows us to say in what way norms such as (1)–(9) are central to the virtues associated with them: in each case, possessing the virtue requires recognizing as a presumptive reason the consideration picked out in the statement of the associated norm. But also, finally, it allows us to answer our question about what explains why some contents of ‘content-neutral’ descriptive considerations are bad, and thus to make a contribution to the intelligibility of value. The norms we have mentioned set out (some of) the considerations a good person should recognize as presumptive reasons for action. But they also tell us what is presumptively bad. It is presumptively bad to act in ways contrary to these norms—to harm people (3), to let them down (5), to coerce them (6), and so on. Again, this is only *presumptively* bad: sometimes we *ought* to coerce people to abandon evil ends. But this is enough to give us an account of when a consideration that supplies us with a presumptive reason fails to give us a good reason. Such a consideration fails to give us a good reason when its content is bad. And its content is bad when the ‘normative orientation’ of that content is contrary to a norm that a good person should recognize. Thus our list of norms gives us not only a list of presumptive reasons, but a list of explanations of when their status as normative reasons is undermined. The problem with malicious enjoyment is that the normative orientation of malice is contrary to (3); the problem with the enjoyment of domination is that it is contrary to (6); and the problem with

vandalistic enjoyment is that it is contrary to:

- (15) X is precious
 ϕ -ing would damage X
 \downarrow
 not ϕ -ing.

The norms associated with these presumptive reasons give the normative orientations that a good person should have. But sometimes, the considerations that give us presumptive reasons can have contents that themselves violate those normative orientations. And when they do, it makes sense for a good person to think that the presumptive reason is only presumptively a reason: its status as a good reason for action has been undermined.²⁷

REFERENCES

- Audi, Robert. 1997. 'Moral Judgement and Reasons for Action', in Garrett Cullity and Berys Gaut (eds), *Ethics and Practical Reason*, Oxford: Clarendon Press, 125–59.
- Crisp, Roger. 2000. 'Particularizing Particularism', in Hooker and Little (2000), 23–47.
- Dancy, Jonathan. 1993. *Moral Reasons*, Oxford: Blackwell.
- Darwall, Stephen. 1983. *Impartial Reason*, Ithaca: Cornell University Press.
- Gibbard, Allan. 1990. *Wise Choices, Apt Feelings*, Oxford: Clarendon Press.
- Hooker, Brad, and Little, Margaret Olivia (eds). 2000. *Moral Particularism*, Oxford: Clarendon Press.
- Jackson, Frank, Pettit, Philip, and Smith, Michael. 2000. 'Ethical Particularism and Patterns', in Hooker and Little (2000), 79–99.
- Korsgaard, Christine M. 1996. 'From Duty and for the Sake of the Noble: Kant and Aristotle on Morally Good Action', in Stephen Engstrom and Jennifer Whiting (eds), *Aristotle, Kant and the Stoics: Rethinking Happiness and Duty*, Cambridge: Cambridge University Press, 203–36.
- Little, Margaret Olivia. 2000. 'Moral Generalities Revisited', in Hooker and Little (2000), 276–304.
- McDowell, John. 1979. 'Virtue and Reason', *Monist* 63, 331–50.
- McDowell, John. 1981. 'Non-Cognitivism and Rule-Following', in Stephen H. Holtzman and Christopher M. Leich (eds), *Wittgenstein: To Follow a Rule*, London: Routledge & Kegan Paul, 141–62.
- McDowell, John. 1995. 'Might There Be External Reasons?', in J.E.J. Altham and Ross Harrison (eds), *World, Mind and Ethics: Essays on the Ethical Philosophy of Bernard Williams*, Cambridge: Cambridge University Press, 68–85.
- McNaughton, David. 1988. *Moral Vision* Oxford: Blackwell.
- McNaughton, David, and Rawling, Piers. 2000. 'Unprincipled Ethics', in Hooker and Little (2000), 256–75.

27. Thanks to an audience at the Australian National University for very helpful comments on an earlier version of this paper.

- Nussbaum, Martha. 2000. 'Why Practice Needs Ethical Theory: Particularism, Principle, and Bad Behaviour', in Hooker and Little (2000), 227–55.
- O'Neill, Onora. 1992. 'Vindicating Reason', in Paul Guyer (ed), *The Cambridge Companion to Kant*, Cambridge: Cambridge University Press, 280–308.
- Parfit, Derek. 1997. 'Reasons and Motivation', *Proceedings of the Aristotelian Society Supplementary Volume*, 99–130.
- Raz, Joseph. 2000. 'The Truth in Particularism', in Hooker and Little (2000), 48–78.
- Stocker, Michael. 1979. 'Desiring the Bad', *Journal of Philosophy* 76, 738–53.
- Velleman, J. David. 1992. 'The Guise of the Good', *Nous* 26, 3–26.

PARTICULARISM AND MORAL THEORY

by Garrett Cullity and Richard Holton

II—Richard Holton

PRINCIPLES AND PARTICULARISMS

ABSTRACT Should particularists about ethics claim that moral principles are never true? Or should they rather claim that any finite set of principles will not be sufficient to capture ethics? This paper explores and defends the possibility of embracing the second of these claims whilst rejecting the first, a position termed 'principled particularism'. The main argument that particularists present for their position—the argument that holds that any moral conclusion can be superseded by further considerations—is quite compatible with principled particularism; indeed, it is compatible with the idea that every true moral conclusion can be shown to follow deductively from a finite set of premises. Whilst it is true that these premises must contain implicit *ceteris paribus* clauses, this does not render the arguments trivial. On the contrary, they can do important work in justifying moral conclusions. Finally the approach is briefly applied to the related field of jurisprudence.

One thing has become clear from recent discussions: moral particularism is not a single doctrine, but a family of doctrines. Garrett Cullity, in his contribution to this symposium, distinguishes particularism about principles from particularism about reasons; and he argues for a modest version of the latter. I want to attend to the other side of his distinction: to particularism about principles. The main argument that particularists have advanced still leaves space for a considerable role for principles; and so leaves open the possibility that principles can play an important role in justifying moral verdicts. My aim is to develop the outline of an account in which they can play such a role. It transpires that if principles are to feature in deductive arguments, they must contain implicit *ceteris paribus* clauses. I try to formulate these with sufficient precision to show that they need not lead to triviality. The result is, I think, a position which, whilst still distinctively particularist, escapes the main charges that have been raised against the doctrine.

I

Characterizing Particularism about Principles. What should particularists about principles (hereafter: particularists) say? Should they deny that moral principles are *ever* true? That is probably the standard interpretation, and it is the one taken by Cullity. On Cullity's account even the weakest form of the doctrine involves denying that there are *any* correct general principles of the form:

Under descriptive conditions D the correct moral verdict is V.

But there is an alternative way of characterizing the particularists' position. Rather than denying that principles are ever true, they could be denying that they are *could ever* be *sufficient*: denying that any set of principles can serve to *capture* ethics. On this second interpretation the particularists' claim is that there is no one set of principles that can be used to determine the correct moral verdict in any situation.

Which of the two interpretations is the better? It seems to me that Cullity's cannot be quite right, for particularists typically endorse the supervenience of the moral on the descriptive. That is, they accept the highly plausible thesis that any two situations that are identical in their descriptive properties will be identical in their moral properties.¹ But then if D were a complete description of a situation (including the claim that it is a complete description!), and if V were the correct verdict in that situation, then the corresponding principle would be universally true: *when-ever* D were realized, V would be the correct verdict.

We can imagine various ways of amending Cullity's account to deal with this worry; indeed, Cullity himself suggests one when he says that, according to particularism, there are no 'exceptionless' principles. This seems to me to be along the right lines. One way of making it precise will become clear once we get clearer on the second, more promising, interpretation of particularism. So let us turn to that.

The second interpretation of the particularist approach involves denying that there is any set of principles which serves

1. I skate over the issue of whether the supervenience is inter- or intra-world. For discussion, see Frank Jackson, Philip Pettit and Michael Smith 'Ethical Particularism and Patterns', in B. Hooker and M. Little (eds) *Moral Particularism* (Oxford: Clarendon Press, 2000), pp. 79–99 at p. 84.

to capture ethics. More exactly, on this second interpretation let us take particularism to be the thesis that

There is no finite set of finite principles that serves to axiomatize ethical evaluation: that is, no finite set of finite principles, such that, given any action fully described in non-moral terms, the principles and the description entail a given moral verdict if and only if it is true.²

The need for talk of a *finite* set of principles comes from the fact that, as we have seen, particularists accept the supervenience of the moral on the descriptive. Given this, it would be possible to give a consistent axiomatization of ethics by simply listing every possible action and the correct verdict. In this sense then, particularists must accept that there is a consistent axiomatization of ethics, at least in a loose sense of 'axiomatization'.³ But since there are infinitely many possible actions, such a list would have to be infinite. What the particularist claims is that there is no *finite* set of principles that serves to axiomatize ethics. Or, more precisely, since a finite list can be packed into a single principle provided that the principle is infinitely long, particularists will have to insist that their claim is that there is no finite set of *finite* principles.⁴

Formulated in this way, particularism about principles remains a radical doctrine. It is clearly incompatible with standard 'principled' theories: Kantianism and Utilitarianism for example. It

2 This way of understanding particularism was suggested by the very similar characterization given in Jackson, Pettit and Smith *op. cit.* Two points of clarification: (i) I take it that we should interpret 'entail' in a basically syntactic sense: the particularist is denying that there is any set of principles from which we can *deduce* which action is right; the principled theorist, in contrast, wants something approaching a decision procedure; (ii) whilst this thesis denies the possibility of deducing whether an action is right given its characterization in *non-moral* terms, there is a parallel but stronger thesis denying the possibility of deducing whether an action is right given its characterization in terms of *thick ethical concepts*. What I have to say in this paper about the former thesis can be easily transferred to the stronger thesis.

3. Loose, since all that we mean by 'axiomatization' is a set of principles that, together with the descriptive facts, entail all and only the ethical truths. Standardly in logic the term is only used if that set is decidable. In this stricter sense particularists presumably deny that there is an axiomatization of ethics. (The need for talk of a *consistent* axiomatization comes from the requirement that the principles entail all *and only* the ethical truths. An inconsistent axiomatization will, of course, entail all the ethical truths; but it will entail all the falsehoods as well.)

4. Although I have expressed this in terms of finitude, the particularists' worries will still be telling if the range of possible actions is not infinite but simply unsurveyably large.

holds that any attempt to discover the principles of morality will be unsuccessful. Nevertheless, there is a sense in which it should not be so surprising. For we know from Gödel's work that a similar claim is true about first order arithmetic:

There is no finite set of finite principles that serves to axiomatize first order arithmetic: that is, no finite set of finite principles, such that, given any sentence in the language of arithmetic, the principles entail that sentence if and only if it is true.⁵

This parallel might give pause to those who argue that particularism about principles cannot be right for ethics, since it would make ethics arbitrary or unlearnable. We do not think that Gödel's results have shown arithmetic to be so.⁶ But rather than pursuing that issue here, I want to follow up another point that is suggested by the analogy.

The unaxiomatizability of arithmetic entails that there is no *one* finite set of axioms that entails all of the truths of arithmetic. It doesn't entail that there are truths of arithmetic which do not follow from any set of axioms (and hence are not formally provable). It is trivially true that any sentence can be proved from some set of axioms, provided we are liberal enough in what will count as an axiom: simply add the sentence itself to the axioms. Indeed, the unaxiomatizability result doesn't obviously rule out the possibility that any truth of arithmetic can be given an *interesting* proof; although here we obviously need to make precise just what it is for a proof to be interesting.

A parallel for the particularist view of ethics should be evident. Interpreted in the second way, particularists are committed to thinking that there is no one set of true principles that entails, and hence justifies, each true moral verdict. But they are not thereby committed to thinking that there are true moral verdicts

5. Second order arithmetic is axiomatizable, but only if 'entail' is read semantically. There is still no finite set of axioms from which the theorems of second order arithmetic can be deduced.

6. Jackson, Pettit and Smith make the unlearnability point. Of course, there is much that we would need to get clear on if we were to make the parallel stick. Let me just mention one point at which it might seem to break down. Gödel has shown that an axiomatization of arithmetic will be, at best, partial; whereas it might be thought that particularists think that any attempted axiomatization of ethics will result in falsehoods. In fact I think that there is no disanalogy here, since I deny this latter claim: a point to which I return (see n. 15 below).

that are not entailed by some true moral principles. That is, whilst they reject

- A** $[\exists X: X \text{ is a finite set of true moral principles}][\forall y: y \text{ is a moral verdict}] y \text{ is entailed by } X \text{ and the non-moral truths}$

they can, consistently with this, accept

- B** $[\forall y: y \text{ is a moral verdict}][\exists X: X \text{ is a finite set of true moral principles}] y \text{ is entailed by } X \text{ and the non-moral truths.}$ ⁷

Let's call the position that involves rejecting **A** but embracing **B** *principled particularism*. The idea is that different moral verdicts will be entailed by different sets of principles; but there is no one set that will entail them all. This is the position that I shall explore in this paper. It might seem that it will involve denying outright the thesis that Cullity takes to be definitive of particularism, the thesis that there are no true moral principles. However, as we shall see shortly, there will turn out to be something right about his approach. Getting to that point will involve getting clear on just how principled particularism can be made to work.

II

Developing Principled Particularism. Whilst particularist writings are not clear on the matter, I rather doubt that anyone has so far embraced principled particularism.⁸ Yet it seems to me to be an interesting and attractive position. To see why it is attractive

7. Again I am assuming that entailment should be read syntactically. Suppose it were read semantically: should particularists still reject **A**? They might contemplate accepting it, on the grounds that this would somehow fix the right interpretation. (Compare the discussion for arithmetic in Vann McGee, 'How We Learn Mathematical Language', *Philosophical Review* 106 (1997), pp. 35–68.) But I have no idea how this would go; in particular, I have no idea what would serve as the relevant axioms.

8. For some particularist discussions of the proper role of principles see Jonathan Dancy, *Moral Reasons* (Oxford: Blackwell, 1993), pp. 66–71; and David McNaughton, *Moral Vision* (Oxford: Blackwell, 1988), Ch. 13. Note that the idea that every moral truth is justifiable from some set of true moral principles is compatible with the idea that there can be regret when a principle is justifiably violated; and with the idea that no non-moral features have a constant 'moral valence' i.e. their presence will always either make an act better, or make it worse. On this last point see below, n. 12.

consider the two roles that principles have often been expected to play. In the first place they might be used in an *investigative* role, enabling us to discover the right thing to do. In the second they might be used in a *justificatory* role: to show that (and perhaps why) certain actions are right, to convince ourselves and others of their rightness.⁹ It is a criticism that is often levelled at particularists that they cannot account for either of these roles.

The advocate of principled particularism will indeed think that principles have a very limited function in an investigative role. Principles will not provide us with anything like a decision procedure for telling whether an action is right. The reason is plain: it will be no good trying to discover whether a moral verdict is true by seeing whether it follows from a certain set of principles, since we can never be sure which set of principles to start with.¹⁰ Nevertheless, the advocate of principled particularism will hold that principles play an important justificatory role. We can show that a given action is right by showing that it follows from some true principles and the non-moral facts. The principled particularist will thus be able to explain why principles play such an important role in our moral lives, whilst at the same time explaining the overblown claims that Kantians and utilitarians have made. Impressed by the fact that we use principles to justify moral verdicts, they have sought to codify the principles that we use; and that cannot be done.

That is why principled particularism is attractive. But there is a problem in developing it, one which comes from the very argument that particularists use to support their position. That argument, which we can call *the supersession argument* runs as follows: given any action whose features are described in non-moral terms, and a principle that says that a action having those features will be good, we can always think of some further feature which is such that, were the action to have that feature too,

9. I take this terminology from Russ Shafer-Ladau, 'Moral Rules' *Ethics* 107 (1997), pp. 584–611. I think that some of his conclusions need to be modified in the light of the distinction made here.

10. Given a strict understanding of what a decision procedure is, it is very unlikely anyway that we should have one for ethics: for it seems likely that to formulate ethical arguments we will need at least first order logic, and we know that we don't have a general decision procedure for validity there. But in a looser sense we might say that we have a decision procedure if we know that all of the sound moral arguments are those which invoke a certain finite set of moral principles. At least then we would know where to start looking for the proof.

it would become a bad action. So the principle is inadequate as it stands. It must be modified, or a further principle must be given outlining the exception. But once we have done that an exception to the amended principle(s) will be found, and we will be forced to amend again. And so on.¹¹

The supersession argument is an argument for denying **A**, i.e. for thinking that there is no finite set of true finite principles that, together with the non-moral facts, entails, for all actions, whether or not they are right.¹² But the supersession argument seems to work equally well against **B**: it seems to follow that no deductively valid argument can be found that takes us from true principles and non-moral facts to the conclusion that any particular action is right. For deductive logic is monotonic: if an argument is deductively valid you cannot make it invalid by adding a further premise. So if the addition of a further premise can make an apparently valid argument invalid, it seems that the argument cannot really have been valid in the first place.

This is a formidable problem, but we should not despair. There is, I believe, a remedy. Lest it seem arbitrary or *ad hoc*, let me

11. Cullity refers to this process as 'undermining' rather than supersession; I prefer the latter term, since it captures the idea that the superseding moral consideration involves a positive, as well as a negative, claim.

12. Let me reinforce a point made by Cullity. The argument for particularism is sometimes said to be holism, i.e. the view that non-moral features do not have a constant moral valence, but are good or bad depending upon the other non-moral features present. But that in itself doesn't entail particularism about principles. A homely example will make the point. Worm based compost heaps work best at a certain level of acidity. Make them too acid or too alkaline, and the worms that eat the rubbish will fail to flourish and ultimately will die. Suppose I have some material that I am thinking of adding to my heap. Will it be good for the heap, or bad for it? I won't know that by knowing only the acidity of the material. Suppose that it is strongly alkaline. If the heap is too acidic, as they often are, then its addition will be good. But if the heap is already strongly alkaline, perhaps as the result of an over-zealous application of lime, then it will be bad. Similarly, the *amount* of material matters: the worms need to be fed, but they don't want to be swamped, or aerobic decomposition might set in and fry them. Compost heap management is thus a holistic business, in Dancy's sense: addition of the same stuff can be good or bad depending on the state of the heap. Nevertheless, it would be ridiculous to say (or at least, ridiculous to say on these grounds) that there is no finite set of principles governing it. We can perfectly well work out what the ideal conditions for the worms are, and derive our principles from that. It is just that the principles must make reference *both* to the nature of the material to be added *and* to the state of the heap: whether it is good to add some material will be a function *both* of the acidity and volume of the material *and* the acidity, volume and worm population of the heap. The argument for particularism is thus not just that moral considerations are holistic, but, in addition, that there are infinitely many of them.

first try to motivate it intuitively. In defending **B** we want to say that certain features of the world, together with certain principles, make a certain action right. The worry then is that there *could be* other features of the world which, together with other principles, which would undermine that verdict by making the action not right. But at that point we want to say something like this:

So what? Why be worried by hypotheticals? If there *were* these other features they would make the action not right. But there aren't. We are concerned with the features that actually do obtain, and they, together with the principles, make the action right.

How do these considerations translate into the availability of a deductive argument? The basic idea is clear enough: we want our deductive argument to state that the facts we're talking about are all the relevant facts. Of course we might be wrong about that claim; there might be other facts that we should have considered. But with just about any argument we might be wrong about the premises. What we want is a deductively valid argument which, if we are right about the facts and about them being all the relevant facts, will take us to the conclusion of what we ought to do.

If the idea is clear enough, showing quite how to implement it is a bit more tricky. To keep things simple, we will consider only moral arguments that contain a single conditional principle with a single universal quantifier and that employ *modus ponens*.¹³ Moreover, the examples I shall present are utterly uninteresting. It is the *form* of the argument that I want to explore.

Let us first introduce some terminology. Suppose we have a set of non-moral predicates $\{F_1, F_2 \dots F_m\}$; and suppose that these occur in a moral principle of the form $\forall x ((F_1x \& F_2x \& \dots \& F_mx) \rightarrow F_cx)$ and in a corresponding set of non-moral sentences $\{F_1a, F_2a, \dots F_ma\}$. Then we say that that principle and those non-moral sentences are *superseded* by another moral principle $\forall x ((G_1x \& G_2x \& \dots \& G_nx \rightarrow G_cx))$ and corresponding set of non-moral sentences $\{G_1a, G_2a, \dots G_na\}$ just in case:

- (i) $(G_1x \& G_2x \& \dots \& G_nx)$ entails $(F_1x \& F_2x \& \dots \& F_mx)$, but not *vice versa*;

13. I don't think that this is too unrealistic an assumption; and it would, I think, be easy enough to generalize, but at the cost of making my presentation very hard to follow.

(ii) $F_c x$ is incompatible with $G_c x$.

For instance, the principle 'If something is a killing, you shouldn't do it' and the non-moral sentence 'This is a killing' is superseded by the principle 'If something is a killing and done in self-defence, you may do it' and the non-moral sentences 'This is a killing' and 'This is done in self-defence.'

The particularists' argument was that for every true principle and set of facts, there was some other true principle and set of *possible* facts that would supersede it. But that does not imply that each true principle and set of facts is *actually* superseded; that is, it does not imply that there is a true principle and a set of *true* non-moral sentences that supersede it. We now want to build into each moral argument the claim that it is not actually superseded. How do we do it? First, we need to add to each argument a new premise, which we might call '*That's it*':

That's it: There are no further relevant moral principles and non-moral facts; i.e. there is no true moral principle and set of *true* non-moral sentences which supersede those which appear in this argument.

And then we need to add a clause in each principle to the same effect. We can do this by adding it as a further conjunct of the antecedent. So we get moral arguments like this

- I. P1 This is a killing
- P2 $\forall x ((x \text{ is a killing} \ \& \ \textit{That's it}) \rightarrow \text{you shouldn't do } x)$
- P3 *That's it*
- C You shouldn't do this.

The particularist claims that this argument is bound to be superseded by other valid arguments, for instance

- II. P1 This is a killing
- P2 This is done in self defence
- P3 $\forall x ((x \text{ is a killing} \ \& \ x \text{ is done in self defence} \ \& \ \textit{That's it}) \rightarrow \text{you may do } x)$
- P4 *That's it*
- C You may do this.

But the fact that II is a *valid* argument that supersedes I does not show that there is a *sound* argument that supersedes I. If there is a sound argument that supersedes I, then I cannot be sound,

since the *That's it* premise in **I** will be false. Thus we can embrace **B**, or more accurately (and assuming that every argument can be phrased as one with a single principle), **B***

B* $[\forall y: y \text{ is a moral verdict}][\exists x: x \text{ is a true, finite moral principle}] y \text{ is entailed by } x, \text{ the relevant non-moral truths (i.e. those that interact with the principle } x) \text{ and } \textit{That's it}.$

Let us stop for a moment and fix up one sloppiness that has been allowed to go unchecked. Earlier I gave a definition of what it is for one principle and set of non-moral sentences to supersede another. But this definition was restricted to sentences of a certain form, and that form made no space for the inclusion of *That's it* in the principle. So we need to add to that definition as follows:

A moral principle of the form $\forall x ((F_1x \ \& \ F_2x \ \& \ \dots \ \& \ F_mx \ \& \ \textit{That's it}) \rightarrow F_cx)$ and a corresponding set of non-moral sentences $\{F_{1a}, F_{2a}, \dots F_{ma}\}$ are *superseded* by another moral principle $\forall x ((G_1x \ \& \ G_2x \ \& \ \dots \ \& \ G_nx \ \& \ \textit{That's it}) \rightarrow G_cx)$ and corresponding set of non-moral sentences $\{G_{1a}, G_{2a}, \dots G_{na}\}$ just in case:

- (i) $(G_1x \ \& \ G_2x \ \& \ \dots \ \& \ G_nx)$ entails $(F_1x \ \& \ F_2x \ \& \ \dots \ \& \ F_mx)$, but not *vice versa*;
- (ii) F_cx is incompatible with G_cx .

This reveals a certain circularity in our definitions: *That's it* is defined partly in terms of the idea of being superseded; and being superseded is defined partly in terms of *That's it*. This is a tricky area, but I don't think that the circularity is pernicious. It is simply that we need to understand the two notions together.¹⁴ Nevertheless, the use of *That's it* might well raise a suspicion of triviality. In the next section I investigate whether this suspicion has any foundation. At this point let us just note that we can finally say what was right about Cullity's characterization of particularism about principles: whilst it need not deny that there are

14. On the general issue see Stephen Yablo, 'Definitions, Consistent and Inconsistent', *Philosophical Studies* 72, (1993), pp. 147–75.

any true moral principles, it does deny that there are any that do not contain *That's it*.¹⁵

III

Is 'That's It' Legitimate? *That's it* involves a degree of self-reference; it contains the expression 'this argument' which refers to the very argument in which it occurs. This has the consequence that the *That's it* premise will have a different content in each argument in which it appears.¹⁶ The same goes for the *That's it* clause in each principle. In turn this means that principles will only have a truth value in the context of an argument. But self-reference isn't bad by itself; and I can't see that what we have here will lead to paradox or contradiction. If there are worries they come either from the thought that *That's it* makes moral arguments trivial, or that it imports unacceptable moral considerations into the premises. Let us take these worries in turn.

Arguments can be trivial in many different ways. One way is to be question-begging. Thus an argument that contained among its premises the claim that a certain action was the right one wouldn't be very useful. But adding *That's it* isn't going to be tantamount to doing that. *That's it* makes the very different claim that no considerations beyond those mentioned in the argument are going to be relevant. And in saying this it is going to be highly contentious. Those disagreeing with a moral argument will frequently contend that the problem is exactly that there are further relevant non-moral facts that are being ignored. Indeed, the *That's it* premise does bring a degree of scepticism to any moral argument. If there are infinitely many principles, and

15. We are also now in a position to answer the worry raised in n. 6, namely that arithmetic and ethics are not analogous, since whilst Gödel showed that any attempt to axiomatize arithmetic would be *incomplete*, the particularist holds that any attempt to capture ethics in a set of principles will actually lead to *falsehood*. Certainly the particularist will hold that many attempts to capture ethics with principles will lead to error; utilitarianism is one example. But not all sets of principles will do so. A good set of moral principles, formulated with *That's it* clauses and employed in arguments containing *That's it* premises will lead to no falsehoods. They will, however, be incomplete, for there will be many circumstances in which *That's it* will not be true. There is thus no important disanalogy with arithmetic here.

16. Nonetheless, we can still think of entailment as defined syntactically. We just need to ensure that within any one argument each occurrence of the indexical expression 'this argument' gets the same interpretation.

infinitely many non-moral facts, then our confidence in the truth of *That's it* cannot amount to certainty. But scepticism of this kind seems quite right.

A second triviality worry concerns whether *That's it* guarantees that any argument that contains it will be valid. Certainly there are self-referential premises that have this character. Thus consider

Valid: This argument is valid.

Valid does seem to bring validity to any argument to which it is added as a premise, at least in an informal semantic sense of validity.¹⁷ Does *That's it* have the same effect? I see no reason to think that it does. It certainly seems that we can give invalid arguments containing it, whether validity is understood syntactically or semantically. For instance

- P1 This is a killing
- P2 $\forall x ((x \text{ is a killing} \ \& \ \textit{That's it}) \rightarrow \text{you shouldn't do } x)$
- P3 *That's it*
- C You should do this.¹⁸

Nevertheless, there remains a worry about triviality.¹⁹ The worry is that every moral principle will turn out to be true; or more precisely, true in every argument in which it occurs (recall that owing to the presence of *That's it*, moral principles only get

17. I.e. the sense in which an argument is valid iff every world in which its premises are true is a world in which its conclusion is true. Proof: suppose, for reductio, that there were an argument A that contained *Valid* as a premise and was invalid. An argument which is invalid is necessarily invalid. So, in every possible world, *Valid* would be false. So there would be no possible world in which the premises of A are all true; so there would be no world in which they are true and the conclusion is false. So A would be valid. Worse still: the argument

P: This argument is valid
C: God exists

seems to be both valid and sound. We have already seen that, since it contains *Valid* as a premise, it is valid; and since its only premise says that it is valid, it must be sound. But then the conclusion will follow, whatever it might be. Clearly there is something wrong with *Valid*. For discussion see Stephen Read, 'Self-Reference and Validity', *Synthese* 42 (1979), pp. 265–74.

18. Note: the problem of an argument with just *That's it* as premise doesn't arise, since the notion of supersession simply isn't defined for such arguments.

19. Thanks to Tim Williamson for raising it.

a truth value when they occur in arguments). Consider one of the principles mentioned above:

$\forall x ((x \text{ is a killing} \ \& \ \textit{That's it}) \rightarrow \text{you shouldn't do } x).$

Let us suppose that that principle is true in whichever argument it occurs. Then presumably the same principle with the contradictory conclusion

$\forall x ((x \text{ is a killing} \ \& \ \textit{That's it}) \rightarrow \text{you should do } x)$

will be false in whichever argument it occurs. Let us take one such (valid) argument:

- (1) P1 This is a killing
- P2 $\forall x ((x \text{ is a killing} \ \& \ \textit{That's it}) \rightarrow \text{you should do } x)$
- P3 *That's it*
- C You should do this.

The principle P2 is a universal generalization, so for it to be false in this argument, it must have at least one instantiation of the form

$((A \text{ is a killing} \ \& \ \textit{That's it}) \rightarrow \text{you should do } A)$

which has a true antecedent and a false conclusion. For the antecedent to be true, both of its conjuncts must be true; and that means that *That's it* must be true. *That's it* says, of course, that the argument (1) is not superseded. Yet, and here is the triviality worry, it looks as though every argument might be trivially superseded. For consider the argument

- (2) P1 This is a killing
- P2 $\forall x ((x \text{ is a killing} \ \& \ \text{Grass is green} \ \& \ \textit{That's it}) \rightarrow \text{you shouldn't do } x)$
- P3 Grass is green
- P4 *That's it*
- C You shouldn't do this.

Here we have simply taken a moral principle that, by hypothesis, we are taking to be true, and a sound argument in which it features; we have inserted into the antecedent of the principle a true non-moral sentence; and we have added that sentence to the argument as a further premise. The resulting argument is sound.

By the definition of supersession that we gave earlier, this argument does seem to supersede the unsound argument, (1), with which we started. So the *That's it* premise that (1) contained is false; and so the moral principle which it contained, which we wanted to be false, has not been falsified after all. Moreover, it looks as though we could play the same trick with any instantiation of that principle; and indeed, with any instantiation of any apparently false principle. It seems that every moral principle will be true: either substantially true, in virtue of featuring in sound moral arguments; or trivially true, in virtue of this supersession trick.

Now perhaps this is not an utterly disastrous result. After all, we have already noted that the presence of *That's it* means that moral principles only get a truth value in the context of an argument. Once we have embraced this relativity to arguments, perhaps we should not worry about which moral principles are true, but about which feature in sound arguments; we might hope that those that are not trivially true are those that do (or could?) feature in sound arguments. (Argument (1) has not been shown to be sound. Indeed, the very argument we gave for thinking that the moral principle it contained was true traded on the fact that it was not, since it traded on the falsity of the *That's it* premise.) However, that is far from obviously right; and anyway it would be good if we had some way of denying truth to certain moral principles.

I see only one way of doing this, which works by being a bit less liberal about what we count as a moral principle in the first place. Recall that the problem arose because apparently false moral principles were superseded by arguments containing principles that were built from true moral principles with extra non-moral clauses—such as the claim that grass is green—inserted into their antecedents. But must we count these gerrymandered constructions as moral principles? We need to be more restrictive: moral principles are minimally contentful. Weaken a moral principle by adding an unnecessary clause to its antecedent, and what you get is not a moral principle at all. It's not obvious quite how the details of this proposal would be worked out.²⁰ But something like it is quite in line with our intuitive conception of what

20. It has something in common with the objection to weakening found in relevance logic; and in working out the details one might expect to find similar difficulties to those encountered there.

a moral principle should be. Indeed, not only is it needed to block the current worry. Some such qualification seems to be needed if we are to say that moral principles must serve to *explain* moral verdicts.²¹

So much for the worry that *That's it* makes moral arguments trivial. What of the worry that it imports too much into the premises? We have already addressed this to some extent when we asked whether it was question begging. But there is further worry that, without actually being question begging, the presence of *That's it* might somehow undermine the whole point of the arguments in which it occurs. The thought is something like this: the interesting part of the principled particularist's claim was that, for any action that we ought to perform, we could give a deductive argument for performing it that used as premises just principles and non-moral facts. But once we add some further premise, we need to be careful that it isn't of a sort which serves to take away the interest of the claim.

Again, I don't think that there is any real concern here. In effect *That's it* consists of two related claims: a claim that there are no further relevant non-moral facts, and a claim that there are no further relevant principles. We thought it quite legitimate for the premises of a moral argument to consist of non-moral facts and moral principles: we were not, after all, trying to derive an 'ought' from an 'is'. So why should it not also include universal generalizations over non-moral facts and moral principles? Admittedly they make use of the idea of relevance; but this we have defined in terms of the notion of supersession; and there does not seem to be anything odd about that definition. Since the quantification ranges over an infinite domain, we cannot think that our understanding of the universal generalizations comes from surveying that domain. But no one other than a very hard-line verificationist would conclude that we therefore do not really understand them, or that they have no significance. They remain eminently falsifiable.

In fact something like *That's it* will be needed by very many moral theories and theories of practical action when put in

21. I haven't broached the difficult issue of moral explanation here. My aim is the limited one of showing how a particularist could embrace principles. I am not saying anything about the further features that such principles must have if they are to serve as explanations.

deductive form. Suppose you think that the thing to do at a party is to talk to the most influential person. You will not know what to do if all you have is a list of people who are at the party, together with details of how influential they are. In order to know who to talk to, you will also need to know that these are *all* the people at the party. Similarly, and more plausibly, suppose you were a utilitarian. Then you couldn't determine which action to perform on the basis of a list of pleasures and pains caused by possible actions of yours. You would need to know, in addition, that these was *all* the pleasures and pains that each action caused; and that these were *all* your possible actions. Indeed it wouldn't help even if, *per impossible*, you listed absolutely every non-moral fact. You would still need to add the premise that these were all the facts: you would need something analogous to *That's it*. Of course it is easy to describe the non-moral facts in such a way that something like *That's it* is smuggled in; we talk about the *total* utility caused by an action. We should realize though that this is just shorthand for the claim that there is this much *and no more*.

It might seem then that the real difference between the role of *That's it* in principled particularism and the role of similar premises in utilitarian arguments is this: whilst both need a premise that says that there are not further relevant non-moral facts, it is only in principled particularism that we need a premise like *That's it* which also makes a claim about there being no further relevant *principles*. But even here the distinction is not so clear. Certainly a utilitarian argument that a certain action is right need only invoke the utilitarian principle that one should perform the action yielding the greatest utility, together with a list of the total utility resulting from each possible action, and a claim that these are all the actions. But if we are to conclude that this is the *unique* action that is right, we will need to know that the utilitarian principle is the *only* correct one; we need to say that this is the only relevant principle. We wouldn't normally put this in to a utilitarian argument, since it is presupposed; but strictly we need it.

IV

Conclusion, and a Brief Application to the Law. I hope that I have said enough to make principled particularism look plausible, and

to quieten worries about the legitimacy of *That's it*. As I said at the outset, I find principled particularism an interesting doctrine. It acknowledges the particularist point that we can always find an exception to any moral principle; whilst at the same time making good sense of the role that principles play in justifying our moral conclusions. It also, I think, makes good sense of our use of principles in teaching people to be moral. We give principles that apply in most cases, but always with an implicit *That's it* condition attached. An insensitivity to this on the part of the learner is likely to be met, not with admiration for their logical acumen, but with impatience at their over-rigid approach. Perhaps too it will be met with the admonition that *the exception proves the rule*. This suggestive but much misunderstood expression comes from the law, an area to which I now turn. I shall say a little more about it at the end of a very brief discussion of how the principled particularism might apply in legal contexts.

In 'Model of Rules I'²² Ronald Dworkin argues that legal positivism, understood as the thesis that all law derives from a set of socially accepted rules identified by a rule of recognition, cannot acknowledge the role of principles. The term 'principle' has a special sense in Dworkin's paper: roughly, unlike rules, principles do not apply in an all-or-nothing fashion, but rather apply with a certain weight, the weight being sensitive to the case.²³ At the end of his paper Dworkin considers the possibility that positivists might simply acknowledge all he says about the importance of principles, but insist that all law derives from a set of socially accepted rules and principles, together with an assignment of their appropriate weights in different cases. This set might be then regarded as the rule of recognition. To this Dworkin responds

This solution has the attraction of paradox, but of course it is an unconditional surrender. If we simply designate our rule of recognition by the phrase 'the complete set of principles in force' we achieve only the tautology that law is law. If, instead, we tried

22. In *Taking Rights Seriously*, (London: Duckworth 1977).

23. Note too that this is likely to be a case where the particularist claim will need to be the stronger one mentioned in n. 2. Legal principles don't typically allow courts to arrive at legal conclusions from premises stated in non-moral vocabulary; rather, they enable the move from certain evaluative legal and moral notions (thick ethical concepts) to verdictive ones.

actually to list all the principles in force we would fail. They are controversial, their weight is all important, they are numberless, and they shift and change so fast that the start of our list would be obsolete before we reached the middle. Even if we succeeded, we would not have a key for law because there would be nothing left for our key to unlock.²⁴

This is weak. If we did compile such a list it is not true that there would be nothing left for our key to unlock. We would have a formidable tool: a complete account of what the correct legal judgement would be in every case. So we need to turn to Dworkin's reasons for thinking that we could never compile it. Here the problem as Dworkin identifies it seems to be mainly technical: the principles shift so fast we could never pin them down. This invites the response that if only we had enough people on the job, with enough resources, then we could get it done. What we want is an argument for thinking that we could never compile the list, no matter what our resources. And it is here that particularist considerations come in. For isn't the thought that, no matter how long and detailed a list of principles and weights we had, we could always think of new cases which would require new principles, or at least, a change of the weights? Dworkin hints at this when he says that the principles are 'numberless' (I take it he means countably infinite, rather than non-denumerable). But why doesn't he say so explicitly? One reason might be that such a claim would seem to be in conflict with another feature of the law. In giving a judgment on a case a court does seem to give a set of rules and principles, together, perhaps, with some weighting on those principles, from which the verdict is said to follow. And in so far as the judgment counts as a precedent, subsequent courts will be required to distinguish the facts in cases in which they want to rule differently. So it might seem then that, once we acknowledge the role of rules and principles in justifying a given decision, we are forced to accept that the total set of such things could, in principle, be given.

I hope that I have shown in this paper that there is no such pressure. Principled particularism provides us with just the resources we need to bring the two theses together. We can think of legal decisions as always containing implicit *That's it* clauses.

24. Ibid. pp. 43–4.

When a distinction is made in a subsequent judgement, it is typically that clause that is denied. Indeed what better way of understanding the legal maxim, mentioned above in a non legal context, *Exceptio probat regulam in casibus non exceptis*: the exception proves the rule in the cases not excepted?²⁵ When we explicitly identify a case as an exception to a rule, we invoke *That's it*. But by treating it as an exception, rather than using it to show that the rule is wrong, we thereby implicitly reaffirm the rule for the cases in which *That's it* is not triggered.²⁶

25. It is unfortunate that this maxim is commonly given the fallacious and far less interesting interpretation (endorsed by *Brewer's*) in which 'proves' is taken to mean 'tests'. This interpretation makes no sense of the use to which the expression is typically put. In discussing the maxim I make no claim about its standing in current Anglo-American law.

26. Thanks to audiences at the Universities of Edinburgh, Glasgow, Leeds, Stirling and Vermont where I tried out some of these ideas; and to Alexander Bird, David Christensen, Garrett Cullity, Jonathan Dancy, Anthony Duff, Frank Jackson, Rae Langton, Peter Milne, Philip Pettit, Michael Smith and Timothy Williamson.