

Principles and Particularisms

RICHARD HOLTON

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.

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

¹I 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 edd. *Moral Particularism* (Oxford: Clarendon Press 2000) pp. 79–99 at p. 84.

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 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.

²This 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.

³Loose, 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.)

⁴Although 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 unsurvayably large.

⁵Second 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.

⁶Jackson, 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

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 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.

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 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

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).

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', *The 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: Basil Blackwell 1993), pp. 66–71; and David McNaughton *Moral Vision* (Oxford: Basil 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.

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, 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

⁹I 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.

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.

¹²Let 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

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 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*;
- (ii) F_cx is incompatible with G_cx .

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

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.

¹³I 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.

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? Firstly 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
 P2 $\forall x ((x \text{ is a killing} \ \& \ x \text{ is done in self defence} \ \& \ \textit{That's it}) \rightarrow \text{you may do } x)$
 P3 *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 is entailed by x, the relevant non-moral truths (i.e. those that interact with the principle x) and *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 $\forall x ((F_1x \ \& \ F_2x \ \& \ \dots \ \& \ F_nx \ \& \ \textit{That's it}) \rightarrow F_cx)$ and a corresponding set of non-moral sentences $\{F_1a, F_2a, \dots, F_na\}$ 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 any true moral principles, it does deny that there are any that do not contain *That's it*.¹⁵

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 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.

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

¹⁵We 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.

¹⁶Nonetheless, 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.

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 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

¹⁷I.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) 265–74.

¹⁸Note: 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.

¹⁹Thanks to Tim Williamson for raising it.

((A is a killing & *That's it*) → 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 | <i>That's it</i> |
| 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 a moral principle should be. Indeed,

²⁰It 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.

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 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,

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.

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.

CONCLUSION & 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 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

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

²³Note 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.

²⁴*Ibid.* pp. 43-4

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. 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 case 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

Richard Holton
Department of Philosophy
University of Edinburgh

richard.holton@ed.ac.uk

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.