

PRAGMATISM AND TELEOLOGY

Christopher Woodard

Contact Details: Department of Philosophy, University of Nottingham, Nottingham,
NG7 2RD, UK. Email: christopher.woodard@nottingham.ac.uk

Abstract

This paper connects two ideas. The first is that some common responses to ethical views are responses to their degrees of pragmatism, where a view's degree of pragmatism is its sensitivity to ethically relevant changes in the actor's circumstances. I claim that we feel the pull of opposing pro-pragmatic and anti-pragmatic intuitions in certain cases. This suggests a project, of searching for an ethical view capable of doing justice to these opposing intuitions in some way. The second central idea is that a theory of pattern-based reasons looks more promising than the obvious alternatives to fulfil this role, amongst Teleological theories at least. Pattern-based reasons are reasons to perform some action because of the goodness or rightness of a larger pattern of action—such as a pattern that a group could perform or that the actor could perform over an extended period—of which the immediate action is a mere part. Existing theories of such reasons share two features that prevent them explaining the intuitions we wish to explain: they consider only one pattern at a time (they are monist), and they treat patterns as eligible only if the agents concerned are willing to realise them (they accept the Willingness Requirement). But we need not accept these doctrines. Moreover, a theory of pattern-based reasons without them is able to explain the pro- and anti-pragmatic intuitions in an elegant way, and has other attractive features.

Pragmatism and Teleology*

Some ethical views are more principled than others. For example, Kant's ethics is more principled than Hooker's Rule Consequentialism, which is in turn more principled than ordinary Act Consequentialism.

What do we have in mind when we think of the principledness of ethical views? One thought might be to classify views as principled if they contain at least one absolute prohibition, but as unprincipled otherwise. A principled view is prepared to rule some things out, ethically.

This suggestion allows us to make sense of the idea that some views are more principled than others, since we can say that a view is more principled to the extent that it absolutely prohibits a greater number of types of action. But the suggestion is otherwise unappealing, since it fails to account for the fact that Hooker's Rule Consequentialism is somewhat principled, and indeed more principled than Act Consequentialism. Neither of these forms of consequentialism absolutely prohibits any type of action.¹ Moreover, it is unattractive, on reflection, to say that the issue over which principled and unprincipled theories divide is only whether any type of action is *always* prohibited. That tends to focus too much attention on the behaviour of ethical theories in logically possible but very unlikely cases. In fact, we distinguish between principled and unprincipled theories according to their behaviour in more likely cases, as well.

This paper explores an alternative understanding of principledness, according to which an ethical theory's being principled is a matter of its degree of pragmatism, in a sense to be explained shortly (in Section 1). This way of understanding principledness

allows a more convincing account of degrees of principledness, and correctly classifies many theories that contain no absolute prohibitions as somewhat principled. Moreover, it picks out a characteristic of ethical views about which we seem to care, and to which we have interesting intuitive responses. Many of us respond to such views under the sway of two opposing intuitions: one pro-pragmatic, and one anti-pragmatic—or so I'll claim in Section 2.

If this is right it suggests a project. It would be at least interesting to find an ethical theory that somehow does justice to both kinds of intuition about degrees of pragmatism or principledness. In the rest of the paper I'll explore some ways of trying to do this, within a teleological framework. The project, so constrained, is to find a *moderately principled* teleological view. Having rejected versions of the doing/allowing and intending/foreseeing distinctions as appropriate models (Section 3), I'll argue that the property of being moderately principled is not best understood in terms of axiological features of ethical views, either (Section 4).

That will bring me to my favoured candidate, which is an unusual theory of *pattern-based reasons*: that is, a theory according to which there can be a reason to do ϕ because of the goodness of a larger pattern of action, P , of which ϕ is a mere part. In the remainder of the paper (Section 5), I'll explain how we can develop an attractive understanding of moderately principled teleology in terms of pattern-based reasons, provided that we depart from some common assumptions about such reasons. Moreover, this account of pattern-based reasons has some merit that is independent of the intuitions about pragmatism.

1. The pragmatic—principled contrast

We sometimes say that a person, argument, or theory is or is not principled. When we say that a person is principled, we might mean that he or she has a great deal of integrity, perhaps. There is a clear set of values or principles that defines this person's ethical commitments, and which he or she is unwilling to breach. Some people are principled in this sense, and others are not.

Similarly, we might try to distinguish between arguments for or against some policy or action according to whether those arguments are principled. Someone might say, for example, that she does not object to capital punishment in principle, but only in practice, because of the imperfections of real trials. Here she is not classifying herself as principled or unprincipled, but is instead classifying some different arguments for and against capital punishment. Finally, we may describe some ethical view or theory as being principled or unprincipled.² Here, I mean by 'ethical view or theory' a theory about the deontic status of actions, or about the normative reasons for or against actions.³

Thus we can distinguish uses of 'principled' according to the kind of object to which the description is applied. And it is likely that we will mean slightly different things when saying that a person is principled than when saying that an argument is principled, say. Even for a single kind of object, though, such as 'ethical theories', there may be several senses of 'principled' abroad. Here I mean to pick out just one, rather than to capture all the nuances of our usage of the term. However, I shall claim that the sense I pick out is one we in fact care about, rather than just something we could possibly mean. In fact I think it is one of the central senses of 'principled' as applied to theories, and one to which we have strong intuitive responses (as I'll explain shortly).

The sense of principledness I have in mind is a descendant of the suggestion we briefly considered at the start, which was that a theory is principled if and only if it prohibits some kinds of action in every case (in which such actions are possible). This suggestion combines the following two ideas:

- (1) A theory is principled only if it prohibits some actions.
- (2) A theory is principled only if it assigns the same deontic status to certain kinds of action, no matter what the circumstances.

To see the importance of the first idea, consider *Anything Goes*. According to this theory, every kind of act is permissible, no matter what the circumstances. This theory satisfies (2), but it is not what we would call principled. Plausibly, we would not call it principled because our idea of principledness incorporates (1).

Let us agree that a theory is not principled unless it prohibits some actions. *Anything Goes* is unprincipled for this reason. Still, we earlier said that Hooker's Rule Consequentialism is more principled than Act Consequentialism, despite the fact that neither theory absolutely rules out any kind of action.⁴ Hence, we had better not adopt the suggestion that principledness is a matter of prohibiting some kinds of action *no matter what*. What then should we say?

My suggestion is that we relax the idea embodied in (2) in two ways. That idea is that principledness is a matter of assigning the *same deontic status* to certain kinds of action in *every case* in which such actions are possible. We can think of this as a particularly demanding version of a more general idea, which is that principledness is a matter of *robustness* in implications across cases, about deontic status *or reasons*. This suggestion substitutes a general concern with degrees of robustness for the specific concern with sameness across every case, and also allows robustness of

implications with respect to reasons to count alongside robustness of implications with respect to deontic status. Let me explain each of these elements in turn.

If a theory assigns the same deontic status to acts of type *T* no matter what the circumstances, it is *wholly* robust in its implications (about deontic status) with respect to acts of that type. As far as acts of that type go, if they are wrong in one case they are wrong in every case. This is part of what we mean by principledness. But it sets the bar for being principled very high. Most ethical views do not contain many exceptionless prohibitions. And even where there is not an exceptionless prohibition, we can distinguish between theories according to how robust their implications are across cases. Indeed, it is because Rule Consequentialism is more robust than Act Consequentialism that it is more principled. Its implications about the deontic status of actions are not wholly robust, since for any type of action we can imagine changes of circumstance capable of changing the status Rule Consequentialism assigns to a token of that type of action. But its implications are more robust than those of Act Consequentialism, which is highly sensitive to changes in circumstances. So we should understand principledness in terms of the general idea of robustness, rather than in terms of the specific concern with lack of exceptions.

Second, we should broaden our concern to take into account a theory's implications about the actor's *reasons*, as well as its implications about the deontic status of actions. Many changes at the level of reasons do not affect judgements about deontic status. Imagine two theories, which agree exactly about the deontic status of every possible act in every case. Now suppose that one of these theories is more sensitive to changes in the actor's circumstances than the other in its implications about the actor's reasons. These different degrees of sensitivity to changes are invisible so long as we understand principledness only in terms of implications about deontic status,

but it is hard to see why they are irrelevant to principledness. We should allow sensitivity to changes at the level of reasons to count as well.

Thus I shall say that a theory is *pragmatic* to the extent that its implications—about the actor's reasons or about the deontic status of her options—are sensitive to changes in the actor's circumstances.⁵ Pragmatism thus understood is the characteristic at stake in (2); it's just that (2) takes an extreme, and unattractive, view about the degree of pragmatism compatible with being principled, and couples this with an unmotivated lack of concern with behaviour in respect of reasons. If we relax the standard embodied in (2) in the two ways I have suggested we arrive at a more attractive understanding of principledness, according to which it combines the following two ideas:

- (1) A theory is principled only if it prohibits some actions.
- (2*) A theory is less principled to the extent that it is more pragmatic (that is, to the extent that its implications about deontic status or about reasons are sensitive to changes in the actor's circumstances).

Thus understood, being principled is not quite the opposite of being pragmatic. Being very pragmatic is sufficient but not necessary for being unprincipled, as Anything Goes reminds us. Still, once we confine attention to theories that prohibit some actions, the degree of principledness of a theory is largely the inverse of its degree of pragmatism.⁶

This kind of principledness may well not be what we always have in mind when we use the term. But it is one of the things we sometimes have in mind. For example, it is one of the things that people seem to have in mind when they say that Kant's ethics is highly principled. They seem to mean, perhaps amongst other things, that his ethics

rules out certain types of action, such as lying, *no matter what* the circumstances. Equally, one of the things that people seem to have in mind when they say that Act Consequentialism is unprincipled is that its implications about whether acts of a certain type are prohibited or required or optional are highly sensitive to changes in the actor's circumstances.

A merit of this way of understanding principledness is that the earlier suggestion can be seen as one version of it. But by substituting a general concern with robustness for the particular concern with exceptionless prohibitions, and widening our concern to include implications about reasons, we reach an understanding of principledness that is capable of distinguishing amongst the vast majority of ethical views. Most ethical views share in common the property of not containing absolute prohibitions; but they differ significantly, nonetheless, in how principled they are.

2. Two intuitions and a project

A theory's degree of pragmatism is a matter of its behaviour overall, not the stand it takes on some single issue. Act Consequentialism, as I have said, is extremely pragmatic. But there is no single component of Act Consequentialism that is sufficient to explain this. Rather, Act Consequentialist theories tend to be highly pragmatic because of the combined effect of their component doctrines.

This illustrates the point that 'pragmatic' is a different sort of description than, say, 'agent-neutral'. In describing a theory as agent-neutral, we are describing one of its component doctrines. In describing a theory as pragmatic, in contrast, we are describing its overall behaviour: the degree to which its implications vary with changes in the actor's circumstances. We have not yet identified what might make a

theory more or less pragmatic, in terms of its component doctrines. The description is not yet a diagnosis. It is rather like describing a person as jumpy, without yet diagnosing why he is sensitive to the slightest change in his surroundings.

This might make us wonder whether there is any merit in thinking about the degree of pragmatism of ethical theories as such, rather than pressing on to search for some underlying diagnosis of each theory's more or less pragmatic behaviour. However, this sort of scepticism seems wrong. We do seem to respond to degrees of pragmatism as such.⁷ Moreover, we seem to have an intriguingly mixed attitude towards degrees of pragmatism, feeling the pull simultaneously of pro-pragmatic and anti-pragmatic, or pro-principled, intuitions.

In thinking about the appeal of being more or less pragmatic, it is important to remember that we are considering the degree of pragmatism of *theories*, not of *persons*, though we seem to respond to the degree of pragmatism, in a similar sense, of persons too. Often we praise pragmatic persons. They are easy to get on with, and tend to get things done, because they focus clearly on the nature of the circumstances at hand, rather than worrying about what would be appropriate in some different case. They ignore sunk costs and let bygones be bygones. They approach each case with an open mind, not expecting that what works elsewhere need work here too. Yet we sometimes worry when a person is extremely pragmatic, especially if he is a political leader. Such a person has no clear boundaries, in the sense of things he absolutely rules out. We may worry that pragmatists might sometimes be a little too willing to abandon established ways of doing things, and we may seek someone a little more principled.

Though these mixed attitudes to pragmatism in persons are interesting, they are not our topic and we should set them aside. We have reasons for valuing principledness

in persons that do not correspond with reasons for valuing principledness in theories. For one thing, there are important considerations of *reputation* that apply to persons but not theories. For example, having a reputation for being principled may discourage certain sorts of attempted manipulation or coercion. For another thing, we may value principledness in persons because we worry about the flaws of powerful people, and we believe that adherence to certain principles can protect us against some of the bad effects of these flaws. Both of these kinds of consideration may help to explain why we value a certain amount of principledness in political leaders. But they do not give us reason to value principledness in theories.

Once we set aside considerations properly pertaining to persons rather than theories, the appeal of a high degree of pragmatism becomes clearer. No doubt it would be absurd to seek an ethical theory that is maximally pragmatic, in the sense of being sensitive to *every* change in the actor's circumstances. Many changes are simply not ethically relevant. But it is appealing to think that ethical theories should be sensitive to every ethically relevant change, in their implications about the actor's reasons or the deontic status of her actions. Call this *the pro-pragmatic intuition*.

Now, of course, different ethical views will differ over what counts as an ethically relevant change. What counts as relevant will depend on what the view in question counts as good, for example. But I shall try to choose examples where this sort of issue does not arise. In any case, our topic is the further issue, of the extent to which ethical views should be sensitive to ethically relevant changes, whatever these are. I claim that there is a fairly widespread *pro-principled intuition*, to the effect that some such changes should make no difference to the actor's reasons.⁸ This may sound distinctly unintuitive, so let me try to illustrate it with some examples.

First, consider *Sexual Harassment*. In this case your colleagues start to harass you, such that you face a choice between wearing what you would like to wear, thus being subject to harassment, and wearing something you would not otherwise choose to wear, thus escaping harassment. The harassment you would suffer, let us say, is annoying and embarrassing but not physically dangerous. Still, it is a harm of sorts to you.

Faced with this case or ones like it, many have the intuition that you should wear what you like, refusing to dress down because of the threat of harassment. Some might even say that you have no reason to dress down, despite the threat of harassment. Yet the advent of the threat seems to be an ethically relevant change in your circumstances. Thus many have the intuition that the deontic status of your options is unaffected by this relevant change, and some have the intuition that even your underlying reasons are unaffected by it. This is, I take it, a fairly appealing principled intuition.

However, some might doubt whether the views that are insensitive to this change would also classify it as ethically relevant. They might suspect equivocation.

Consider, then, *Bomb Threat*. In this case, a campaign of bombing seems to be underway in your city. There is a credible threat of a further bombing in the coming weeks. As in *Sexual Harassment*, there is a threat of harm if you continue to live as normal; but in this case any harm suffered would be much more severe.

I assume that no plausible ethical theory could treat the advent of this threat as an irrelevant change in circumstances. Yet whenever such a threat exists, many people say ‘we should carry on as normal’. This response is *very* common. What these people seem to be saying is that the advent of the serious threat should make no difference to how they live. Taken this way, this is a startling claim. It seems to

express the view that this change in circumstances, which is admittedly ethically relevant, should be, in a way, ignored.

These people probably don't mean exactly this. Most of us would want to take avoiding action if the threat were severe enough, which means that we are tracking the change. But those who say 'carry on as normal' don't quite mean 'carry on as normal, unless the threat is severe enough', either. What then do they mean?

My hypothesis is that in this case and Sexual Harassment many people are torn by the two opposing intuitions I mentioned. The pro-pragmatic intuition speaks in favour of Act Consequentialism and other very pragmatic views, and against principled views. But opposing it is the anti-pragmatic, or pro-principled, intuition, to the effect that even ethically relevant changes should sometimes be ignored, or in some way not allowed to make a difference. It is this intuition that people are expressing when they say, under threat of bombing, that life should 'carry on as normal'. It seems to be quite powerfully held.⁹

More generally, the pro-principled intuition seems to be engaged by cases in which the actor faces what we might call a *regrettable choice situation*: that is, one in which the first-best outcome is excluded by some obstruction. In Sexual Harassment, for example, the first-best outcome (wearing what one likes without the prospect of harassment) is excluded by the harassers' behaviour. In Bomb Threat, the first-best outcome (living as normal without threat of bombing) is excluded by the bombers' behaviour. In contrast, where the first-best outcome is available to the actor, the pro-principled intuition seems not to be engaged. Note that in both of these cases, *other agents acting wrongly* (or threatening to do so) exclude the first-best outcome. This is not true of all regrettable choice situations, but it does seem to be true of those

cases that most strongly engage the pro-principled intuition [Woodard 2000: 410-418].

If my hypothesis is right, it helps to explain the continuing fascination of debates between consequentialists and deontologists, and between Act and Rule Consequentialists. Anyway, the hypothesis suggests a project, of finding an ethical view that somehow does justice to both intuitions. As I shall put it, it would be interesting to find a moderately principled view. And I shall assume that the most promising place to look for such a view is within the class of teleological views, since these seem best placed to do justice to the pro-pragmatic intuition. (However, I do not want to insist on this point, and I am not against looking for moderately principled nonteleological views.) Thus my project is to look for a *moderately principled teleological view*, by which I mean one that does justice to both the pro-pragmatic and pro-principled intuitions in some way.

3. Abandoning consequence-neutrality

I will understand this search in terms of possible additions to or departures from Act Teleology. Act Teleology is just like Act Consequentialism, except that its axiology need not be agent-neutral. Like Act Consequentialist views, Act Teleological views tend to be extremely pragmatic, and thus they gain support from the pro-pragmatic intuition. My method, then, is to consider possible changes to Act Teleology that would help the resulting view to do justice also to the pro-principled intuition. I will consider three such possible changes, though no doubt others would be worth considering as well.

The first departure is from Act Teleology's *consequence-neutrality*. Standardly, Act Teleology treats every consequence for one and none for more than one. This contributes to its high degree of pragmatism; or at least, some ways of departing from it are sufficient to reach a more principled view. I have in mind views that incorporate some version of the distinction between doing and allowing, or the distinction between intended and merely foreseen consequences.

Such views discount merely foreseen or allowed consequences: in determining the actor's reasons or the deontic status of her actions they count them for less, other things equal, than intended consequences, or consequences brought about in the doing way. The rate of discount may be shallow or steep. As it gets steeper, changes in the actor's circumstances that affect the total consequences of her actions without much affecting what she can do or intend recede in importance. Thus views incorporating one of these distinctions tend to be less pragmatic than Act Teleology, and we can imagine tweaking their degree of principledness by adjusting the rate of discount.

Of course, there are well-known problems in formulating these distinctions, and in explaining why they should have the ethical significance that these views claim for them.¹⁰ I will not dwell on these well-known problems, however. Suppose the distinctions can be drawn well enough, and some rationale for discounting is possible; or at any rate, set these issues aside. I want to point out two other features of all such views.

One is that they focus attention on the *actor's* agency. In effect, they accuse extremely pragmatic views of going wrong by failing to take account correctly of the significance of the actor's own agency. She is related to some consequences of her own actions in a way richer than merely being a cause of them—that is, she is related to some effects of

her acts, but not all, as the agent of them. Versions of the doing/allowing and intending/foreseeing distinctions all seek to flesh out this underlying idea, I take it. They shine a spotlight on the actor's agency, accusing Act Teleology of failing to appreciate its full significance.¹¹

One upshot of this first feature is that views incorporating one of these distinctions tend to treat cases in which the actor is faced with second-best choice because of another agent in the same way as cases in which she is faced with a similar choice because of natural forces. That is, such views tend not to track distinctions between types of regrettable choice situation according to the causes of the situation. So far as doing/allowing and intending/foreseeing are concerned, it doesn't matter much whether the actor faces a choice between killing one herself or seeing five killed because of Pedro, or because of a runaway trolley, or because a river has burst its banks.

However, these distinctions in types of regrettable choice situation do seem to matter intuitively. In particular, the pro-principled intuition seems stronger when the first-best outcome is excluded by others acting wrongly than in other cases. For example, it seems to arise more strongly with respect to Jim and Pedro than it does with respect to a structurally similar Trolley Problem. If that is right, it suggests that any error in extreme pragmatism lies in failing to take adequate account of the fact that *others* in the actor's environment are agents. Thus, to the extent that our pro-principled intuitions track the distinctions in types of regrettable choice situation, we have some reason to doubt that distinguishing between doing and allowing, or intending and foreseeing, can really explain them.

Since this is somewhat speculative, I'll put more weight on the second observation. Appeals to doing/allowing or intending/foreseeing do not explain the sense we have

in these cases that there is a *conflict* between different types of consideration at work. I take it that if you have both the pro-pragmatic and the pro-principled intuition, you experience them as being in tension with each other: as suggesting that there are reasons to be pragmatic and reasons to be principled conflicting with each other in a single case. Recall Bomb Threat. As I said, people who express the pro-principled intuition don't seem to mean that there is no reason to be pragmatic, and take avoiding action. Instead they seem to mean that, opposing this reason, there is another reason for carrying on as normal. The intuitions point to conflicting underlying reasons.

If that is right, appealing to some version of the doing/allowing or intending/foreseeing distinctions cannot be the whole explanation of these intuitions. Discounting consequences in this way does not generate conflicting reasons: it merely suggests that there may be more to be said for doing the 'principled' thing (carrying on as normal, not dressing down, not shooting an Indian) than is implied by Act Teleology, and correspondingly less to be said for doing the 'pragmatic' thing than Act Teleology implies. But this is a conflict between Act Teleology *and these other ethical views*, not a conflict *within the actor's reasons* according to these other ethical views.

Merely amplifying the importance of some consequences as compared with others does not generate the sense of conflicting reasons that is part of the phenomenology of the intuitions I am discussing. Hence I conclude that this sort of departure from Act Teleology will not enable us to do justice to the opposing intuitions about pragmatism.

4. Adjusting axiology

Consider then the second strategy, which is to modify one's axiology. This is perhaps the most obvious way to explain the sense that there are conflicting reasons at work in these cases. Starting with an axiology, V , according to which there is a reason (for example) to dress down, in Sexual Harassment, or to take avoiding action, in Bomb Threat, but no countervailing reason to carry on as normal, we can adopt a modified axiology, V^* , sufficient to explain the countervailing reason. In doing this, we can stick with Act Teleology throughout.

There is no restriction within teleology on what has value. One can attach negative or positive intrinsic value to certain types of action, for example. To deal with the sort of cases I have been discussing, one could appeal to these values. One might even think that all such cases could be dealt with in the same way, by postulating that acts of the type, *adjusting one's actions in the light of expected wrongdoing*, have intrinsic disvalue. This would seem to give a standing explanation of what I have been calling pro-principled intuitions. But even if this blanket approach does not work, one could appeal to different values in different cases. For example, one could explain the reason not to dress down, in Sexual Harassment, by saying that dressing down is an instance of *accepting unfair burdens*, and attributing intrinsic disvalue to acts of this type.

Moreover, one could adjust the degree of robustness of the pro-principled reasons by adjusting the degree of value accorded to the types of act in question. If, say, we attach great disvalue to acts that accept unfair burdens, we might be able to explain a reason against dressing down that is quite robust across cases. And supposing that we do not give infinite disvalue to acts of this type, we will still be able to explain how

there can be exceptional cases. So this looks like a more promising model for a moderately principled teleology.

I do not doubt that adjustments to axiology can explain the sense that there are reasons not to do the pragmatic thing in these cases. The question is whether these adjustments seem likely to generate the best sort of account of the balance of reasons. Though they might be the most promising explanation in some cases, there are at least two reasons to doubt whether adjustments to axiology can be the best general explanation of the clash of pro-pragmatic and pro-principled intuitions.

Brad Hooker has pressed one problem with this style of theorising (though he was not addressing the pro-pragmatic and pro-principled intuitions specifically). As he points out, each time we build some extra complexity in to our axiology we rob the resulting teleological theory of explanatory power. He is careful not to say that we should never do this, but instead that we should prefer not to do it, since,

... other things being equal, a theory that does not posit intrinsic goodness or badness as properties of acts makes fewer assumptions. And if a theory making fewer assumptions can explain just as much as a theory making more assumptions, the theory making fewer assumptions is better. [Hooker 2000: 34].

Of course, this type of consideration does not suggest that we should never seek to account for complex judgements by building extra complexity into our axiology, since sometimes other things will be unequal. But Hooker's point does suggest, at least, that we should be on the lookout for other kinds of explanation, and that we should be aware of the danger of postulating extra axiological complexity too readily.

The second point is a generalisation of a familiar difficulty with modelling side-constraints by making axiological adjustments. What is the problem with side-constraints? It is not that they have to be exceptionless, because they do not. Instead the problem is of explaining why the prospect of preventing more acts of the constrained type is not always sufficient reason to perform one act of that type, where there are not other extraneous considerations. Take a simple case: if killing is constrained because of the badness of killing, why isn't killing *one* allowed in every case where doing so is necessary to avoid *two* killings?

The solution is obvious and well known; it is to individuate types of action more narrowly, so that *my killing* is a different type of act than *your killing*, and *killing now* is a different type of act than *killing later*. We can then attach different values to these different types of action, and succeed in modelling side-constraints [Broome 1991: 3-16]. However, there is a cost to modelling side-constraints this way. The resulting axiology may be hard to defend. Certainly it seems difficult to explain why *my killing* should matter more (even so far as my reasons are concerned) than *your killing*, for example; and the same goes for *killing now* as compared with *killing later*.

As I said, the problem is not one of excluding exceptions. We want to allow for exceptions, and the axiological model of side-constraints certainly allows that. Instead the problem is that, if we want to model the balance of reasons by adjusting our axiology, we may find that the axiology gets distorted. Thus this point dovetails with the previous one. Hooker reminded us of the costs of always making axiological adjustments, while the present point is that there may be good reasons not to do so when dealing with specific kinds of reason.

The familiar problems in modelling side-constraints are just one instance of a more general problem. The general problem arises when we seek to account for any reason that plausibly exists, but seems not to fit in Act Teleology, by making adjustments in the axiology. This is a standing temptation, because Act Teleology is such an appealing form of theory. But we can end up with *ad hoc* or otherwise unconvincing axiologies by doing this [Hooker 2000: 34]. We have to say, for example, that Jim's reason not to shoot any Indians is a result of the greater importance, for his reasons, of his killing as compared with Pedro's killing. This seems to give the wrong account of Jim's reasons. If there is a reason for Jim not to shoot any Indians, it must somehow derive, like his reason to shoot one of them, from the importance of the Indians' lives, or their interests. There are not two kinds of good at stake in this case, but only one. Theories that force us to say that something else matters in this case are for that reason implausible.

5. Pattern-based reasons

I objected to versions of the doing/allowing and intending/foreseeing distinctions (as models of moderately principled teleology) by saying that they failed to account for the sense that there is a conflict between pragmatic and principled reasons in the cases we are considering. That naturally suggested a different possible model, which accounts for the conflict of reasons by adjustments in axiology. However, I objected to this second model by saying that it can lead to distorted or otherwise unconvincing axiologies. One example, I claimed, was Jim's case, in which the method of axiological adjustment drives us to postulate implausibly that there is more than one kind of good at stake.

It may seem that I am demanding the impossible. For the conjunction of my objections is that we wish to find some way of accounting for conflicting reasons even in cases where only one kind of good is at stake. But according to a standard assumption this is impossible.

The assumption is that value monist teleology cannot explain genuine conflicts of reasons. Of course, value monists can explain surface-level conflicts of reasons. If *my enjoyment* is all that matters, we can easily understand a sense in which I have one reason to go to the cinema tonight, and an opposing reason to go for a drink, say. But this is not a deep kind of conflict. I have a decisive reason to do whichever of these things I would enjoy the most. The conflict in my reasons really reflects my lack of certainty about which thing I would enjoy most. According to the standard assumption, though, value monists cannot explain the deep sense of conflict to which Jim seems to be subject—the sort of conflict capable of making sense of regret.

However, the standard assumption about value monism is wrong. In fact, we can account for genuine conflicts of reasons by adopting a different sort of pluralism, which does not distort our axiology; thus my objections to the first two ways of modelling moderately principled teleology do not add up to a demand for the impossible. The key to this different form of pluralism is the idea of pattern-based reasons.

As is well known, Act Teleology is an *act-based* view. This means that it characterises the reasons for or against each of the actor's options, and thus the deontic status of those options, in terms of the goodness of those options themselves. To explain the reasons for or against an act, it points to the goodness or badness of that act, as compared with the goodness or badness of the actor's other options. Put somewhat crudely, the reasons for an act depend on the goodness of that act.

Now consider *Collective Teleology*. According to Collective Teleology, the actor is required to play her part in the best collective action that her group could perform, in the circumstances. Collective Teleology deals in what we might call ‘group-based’ reasons: it claims that the actor’s reasons depend on the goodness or badness of an array of group actions, of which her options are mere parts.¹²

Next consider *Plan Teleology*. According to Plan Teleology, the actor is required to perform each part of the best extended sequence of actions that she could perform, over some period, given her present circumstances. Plan Teleology deals in what we might call ‘plan-based’ reasons: it claims that the actor’s reasons depend on the goodness or badness of an array of extended sequences, or plans, of which her present options are mere parts.

Both group-based and plan-based reasons are species of what we can call ‘pattern-based’ reasons. Such reasons, if they exist, are reasons to perform parts of larger patterns, *because of* the goodness or rightness of those larger patterns. The existence of pattern-based reasons is of course a matter of philosophical controversy.¹³ (It is important to see that there can also be act-based reasons for performing parts of larger patterns. No one disputes that.) However, I want to suggest that a theory of such reasons, with two particular features, is a better model of moderately principled teleology than those we have considered up to now.

The first feature is *pluralism*, though pluralism about patterns rather than the more familiar pluralism about values. A credible theory of pattern-based reasons should allow reasons associated with *several* different patterns to bear on a single deliberative problem. Oddly, though several theories of pattern-based reasons exist, none to my knowledge has adopted this feature. Existing theories are all monist in this dimension.¹⁴ But pluralism about patterns is very helpful in accounting for

typical intuitions, and it is also independently credible. Let me take each of these points in turn.

Pluralism about patterns helps explain typical intuitions in just the sort of case we have been considering. Like pluralism about values, it can explain conflicts between reasons. Consider what we might say about Sexual Harassment. The actor has, perhaps, a group-based reason not to dress down, since this is her part in the best pattern of action that the group (consisting of her and her colleagues) can perform, in respect of dressing for work. The best pattern, let us say, involves each person wearing what he or she likes (perhaps within certain limits), and also not subjecting his or her colleagues to harassment or unwelcome attention because of their choices. The actor has a part to play in this optimal pattern, and it consists of (a) not dressing down, but instead wearing what she likes, within the relevant limits, and (b) not harassing others. Thus the sense that she has a reason not to dress down can be explained by appeal to group-based reasons.

However, it is attractive also to allow that she has act-based reasons in this case. She has reason to consider the goodness or badness of her different options themselves, and the harassment she would receive if she were to wear what she likes is one element of this. What's more, as I've said many times, this fits our intuitions: she does have a pragmatic reason to dress down, and this may even be a decisive reason if the threat of harassment becomes severe enough.

Secondly, pluralism is independently credible. To see this, ask why anyone should think that membership of a favoured pattern can itself provide a reason to perform an action. The answer must be that an action's *membership properties* can support reasons, just as its *causal properties* (and perhaps its intrinsic properties) can. The

membership relation, on this view, is capable of supporting reasons: it makes a reason out of the value of the pattern to which the action is related by membership.¹⁵

Of course, not everyone will accept this, but it is difficult to see how else we could give a rationale for belief in pattern-based reasons, independently of fit with intuitions. But if we do accept that membership properties are sufficient to support reasons, it is incredible to think that causal properties cannot. Thus, anyone who believes in pattern-based reasons ought to believe also in (instrumental) act-based reasons. If that is right, any credible theory of pattern-based reasons will be pluralist in the way I have indicated.

The other important feature consists of rejection of a certain widely held constraint on pattern-based reasons (widely held, that is, amongst theorists of pattern-based reasons). According to this constraint, which I call the Willingness Requirement (WR),

WR There is a pattern-based reason to play one's part ϕ in a pattern P performable by a group G only if the members of G are willing to perform P .¹⁶

(To cover plan-based reasons, allow future selves to form a group G with the actor's present self.) This constraint seems plausible on first sight, but we should reject it.

First note that it blocks the explanation of pro-principled intuitions that I have been canvassing in cases like Sexual Harassment and Bomb Threat. In these cases, the relevant agents are unwilling to cooperate in producing the favoured pattern. Worse, it seems to be just those cases with this feature that give rise most strongly to the pro-principled intuition. So the WR implies that pattern-based reasons do not exist in just those cases where we are most likely to have pro-principled intuitions.

That is insufficient reason to reject the WR. But we have other reasons to reject it, fortunately. Its plausibility stems from two main theoretical functions that it looks well suited to perform; but in fact these functions can be performed as well (in one case) or better (in the other) by other devices.

The first function is *avoiding recklessness*. We might worry that belief in pattern-based reasons unconstrained by the WR spells support for reckless actions. For example, we might worry that it would imply that each driver in the UK should drive on the right hand side of the road, if necessary doing so unilaterally. After all, a world in which every driver in the UK drove on the right would be better than the current one, since there would be fewer accidents when UK drivers drove abroad or when drivers from other countries drove in the UK.¹⁷ But it would be crazy to drive unilaterally on the right.

Pluralism allows us to avoid the reckless implication, however. It allows us to say that each UK driver has a reason to drive on the right hand side, since this is each driver's part in the optimal pattern, but that this is overwhelmingly outweighed by a countervailing act-based reason to drive on the left, where others are expected to drive on the left.

Notice that according to this explanation, facts about others' willingness to cooperate in producing the favoured pattern matter, but just not in the way the WR claims. According to the WR, their unwillingness is sufficient to prevent the existence of a pattern-based reason. According to the explanation I just gave, their unwillingness leaves the pattern-based reason unaffected, but changes the nature and strength of the actor's *act-based* reasons. Because others are reliably expected to drive on the left, the actor has a strong act-based reason to drive on the left. (If they weren't expected to drive on the left, she would have no such reason to do so herself.) Facts

about others' propensity to drive on the left affect the actor's reasons, but they affect her act-based not her pattern-based reasons. Of course, in affecting the nature and strength of her act-based reasons, these facts do affect the *relative* strength of her countervailing pattern-based reasons—but that is quite different from the claim made by the WR.

So, pluralism avoids recklessness just as well as the WR does. The other theoretical function to which the WR seems well suited is *excluding mere agglomerations*. Any credible theory of pattern-based reasons must exclude some possible 'patterns' of action as incapable of supporting pattern-based reasons, because they are arbitrary patterns. It is easy to think of examples. Combine any of my actions today with the first occasion on which the President of France scratches his nose this evening, and you have a mere agglomeration. For example, the 'pattern' <I eat my breakfast; President Chirac scratches his nose for the first time this evening> is a mere agglomeration. It is a pattern in a sense, but no one should think that the fact that eating my breakfast is part of it has anything to do with my reasons.

We might worry that we cannot exclude mere agglomerations such as this if we reject the WR. But in fact we can. Consider an analogy: suppose we are interested in the ordinary act-based reasons I had for and against eating my breakfast this morning. The values at stake, let us say, were my pleasure, my nutrition, and my optimal preparation for a lecture I had to give today. The first two values speak in favour of my eating breakfast, while the third may not. (Perhaps I would have done better, in this respect, to read through my lecture one more time.)

My act-based reasons depend on comparing my various options. What were my options? Well, such things as eating my breakfast, reading through my paper one more time instead, and so on. But these are not all of the things I could have done. I

could also have hummed the Marseillaise, for example; and due to my prowess, I could have combined humming the Marseillaise with eating my breakfast. But we do not count *eating-and-humming* as another of my options, in addition to *eating*, even though it is something I could have done. This is because the humming component is *practically irrelevant*: it bears no relation to the values at stake, and so cannot bear on my reasons, given those values.¹⁸

We all rely on a tacit understanding of practical relevance when considering an actor's options in ordinary act-based reasoning. But the same idea explains why the pattern involving my eating breakfast and President Chirac's scratching his nose is a mere agglomeration: his scratching his nose is practically irrelevant, given the values at stake in considering my reasons. For that reason, any pattern that includes his scratching his nose is a mere agglomeration, as far as my reasons to eat breakfast this morning are concerned.

So we do not need to appeal to the WR to exclude mere agglomerations. What's more, though the WR is superficially attractive for this function, it is not in fact very well suited to it. Consider what it says about the pattern involving Chirac's scratching: this is ineligible to support pattern-based reasons because Chirac is unwilling to participate in it. This brings us to an issue about the formulation of the WR. Must the other agents be willing merely to behave in the specified ways, or also be willing to do so under the description of cooperating in producing this pattern? Presumably Chirac is willing, in the first sense, to scratch his nose this evening, so we had better appeal to the stronger version of the constraint. But then suppose that, for some odd reason, Chirac became willing to participate in the pattern involving his scratching and my eating, under that description. That pattern would *still* be a mere agglomeration, but the WR, even on the stronger interpretation, would not explain why.

Thus, despite initial appearances, the exclusion of mere agglomerations has nothing to do with others' willingness. Instead, it is a matter of practical relevance to the case at hand.

6. Conclusion

A theory of pattern-based reasons that is pluralist and rejects the WR is an attractive model for a moderately principled teleological view. It allows us to explain the pro-pragmatic intuition and the pro-principled intuition: the pro-pragmatic intuition reflects the force of act-based reasons, while the pro-principled intuition reflects the force of pattern-based reasons. For a reason associated with a certain pattern, *P*, in which another agent is involved, changes in that agent's intentions (in respect of the behaviour concerned) have no effect on the reason. Thus, some elements of the overall pluralist view are insensitive to changes in these facts, though the complete view is not, since it also contains an act-based component. In this way, this kind of pluralism explains why we have the sense, in the cases under discussion, that there is a clash of reasons.

The sort of view I have sketched here is not a complete ethical theory, but needs supplementing in various ways: by a theory of the good, and by an account of how the reasons associated with various patterns interact. Possibly, it may need supplementing also by further restrictions on the eligibility of patterns. As it stands, it is merely a form of ethical view, in the same sense that bare Rule Teleology and Act Teleology are forms of view. But that does not make it unworthy of study, since we do not yet know enough about the different possible forms of ethical theory, and their comparative merits and demerits.

In fact, the pluralism I have described has some features in common with Hooker's Rule Consequentialism, though there are also important differences. Hooker's view tacitly rejects the WR, and it is minimally pluralist, combining a general concern with the goodness of extended patterns with an exceptional concern with the goodness of individual actions (governed by the disaster prevention rule).¹⁹ My view retains and motivates the first feature, and extends the second, to reach a more thoroughgoing form of pluralism. The result is a more flexible form of moderately principled teleology, which is better placed to give considerable weight to pragmatic considerations. Moreover, it provides an explicit rationale for caring both about the goodness of individual acts and the goodness of extended patterns of action, which helps it to escape the incoherence objection that is repeatedly made against Rule Consequentialism.²⁰

Moreover, the view I have sketched has other pleasing features. It allows teleologists to explain conflicting reasons without needing to postulate conflicting values. This provides the utilitarian, for example, with a response to the charge that she cannot give an adequate account of regret. And it allows us to explain why Jim is subject to conflicting reasons, even though only one thing matters in his case: the welfare of the Indians. In contrast, other views require us to postulate that something else matters, such as Jim's moral integrity. For that reason, these other views are implausible. They force us to fragment our ethical concerns when, in fact, we do not need to do so.²¹

References

Bacharach, Michael 1999. Interactive team reasoning: A contribution to the theory of co-operation, *Research in Economics*, 53: 117-147.

Bennett, Jonathan 1995. *The Act Itself*, Oxford: Clarendon Press.

Broome, John 1991. *Weighing Goods: Equality, Uncertainty and Time*, Oxford: Basil Blackwell.

Gruzalski, Bart 1986. Parfit's Impact on Utilitarianism, *Ethics*, 96: 760-783.

Hollis, Martin 1998. *Trust within reason*, Cambridge: Cambridge University Press.

Hooker, Brad 2000. *Ideal Code, Real World*, Oxford: Clarendon Press.

Hurley, S. L. 1989. *Natural Reasons*, New York: Oxford University Press.

Hurley, S. L. 1991. Newcomb's problem, prisoners' dilemma, and collective action, *Synthese*, 86: 173-196.

Hurley, S. L. 1994. A new take from Nozick on Newcomb's Problem and Prisoners' Dilemma, *Analysis*, 54: 65-72.

Hurley, S. L. 2004. Rational Agency, Cooperation and Mind Reading, in *Teamwork*, ed. Natalie Gold, London: Palgrave: 200-235.

Jackson, Frank 1997. Which Effects?, in *Reading Parfit*, ed. Jonathan Dancy, Oxford: Blackwell: 42-53.

Kagan, Shelly 1989. *The Limits of Morality*, Oxford: Clarendon Press.

Kant, Immanuel 1996 (1797). On a supposed right to lie from philanthropy, in *Immanuel Kant: Practical Philosophy*, ed. and trans. Mary J. Gregor, Cambridge: Cambridge University Press: 609-615.

McClennen, Edward F. 1990. *Rationality and Dynamic Choice*, Cambridge: Cambridge University Press.

McMahon, Christopher 2001. *Collective Rationality and Collective Reasoning*, Cambridge: Cambridge University Press.

Nozick, Robert 1974. *Anarchy, State, and Utopia*, Oxford: Basil Blackwell.

Nozick, Robert 1993. *The Nature of Rationality*, Princeton, NJ: Princeton University Press.

Otsuka, Michael 1991. The Paradox of Group Beneficence, *Philosophy and Public Affairs*, 20: 132-149.

Regan, Donald 1980. *Utilitarianism and Co-operation*, Oxford: Clarendon Press.

Scheffler, Samuel 2004. Doing and Allowing, *Ethics*, 114: 215-239.

Sugden, Robert 2000. Team Preferences, *Economics and Philosophy*, 16: 175-204.

Tännsjö, Torbjörn 1989. The Morality of Collective Actions, *The Philosophical Quarterly*, 39: 221-228.

Williams, Bernard 1973. A critique of utilitarianism, in *Utilitarianism: for and against*, J.J.C. Smart and Bernard Williams, Cambridge: Cambridge University Press: 75-150.

Woodard, Christopher 2000. The Concept of Acquiescence, *The Journal of Political Philosophy*, 8: 409-432.

Woodard, Christopher 2003. Group-Based Reasons for Action, *Ethical Theory and Moral Practice*, 6: 215-229.

Woodard, Christopher 2007. *Reasons, Patterns, and Cooperation*, New York: Routledge.

* For discussion of these issues over many years I am most grateful to Susan Hurley. For helpful comments on earlier versions of this material I am grateful also to the late Michael Bacharach, and to Nick Bardsley, Krister Bykvist, Garrett Cullity, Nir Eyal, Bill Fish, Guy Fletcher, Brad Hooker, Shelly Kagan, Alastair Norcross, Ben Saunders, Julia Watt, and audiences at Dartmouth College, Nottingham, and Oxford.

¹ By an 'absolute' prohibition, I mean one that is exceptionless. Like other sophisticated Rule Consequentialisms, Hooker's includes a disaster prevention rule, which renders any prohibitions mandated by the other rules non-absolute [Hooker 2000: 99; see also 133-134]. Also, I assume here that we can exclude from the class of prohibitions those that govern gimmicky types of action, such as *failure to maximise the good*.

² The contrast between pragmatic and principled arguments is somewhat independent of the contrast between pragmatic and principled theories, since the former raises issues also of ideal vs. nonideal theory.

³ Other kinds of ethical theory might be about the correct way to deliberate, say, or the responsibility of agents for different actions. Here I focus exclusively on theories about the deontic status of actions—that is, whether they are *required* or *prohibited* or *optional*—or about the normative reasons for or against an action. Agents may or may not be aware of their normative reasons or the deontic status of their actions.

⁴ Indeed, even Nozick is attracted to the idea that violations of side constraints might be permitted in some circumstances [Nozick 1974: 29-30n.]. Yet an account of principledness presumably ought to classify his ethics as highly principled.

⁵ How should we balance increased sensitivity in one of these dimensions against decreased sensitivity in the other, to reach an overall judgement about a theory's degree of pragmatism? I do not know, and shall not address that issue here. For our purposes it is sufficient to have rough and ready intuitive judgements about degrees of pragmatism, and I assume that we do have those.

⁶ Only 'largely', since some theories will prohibit fewer types of action than others, and the former are plausibly regarded as less principled than the latter, *ceteris paribus*. In what follows I shall ignore this dimension of principledness, however, and focus on the dimension corresponding with (lack of) pragmatism.

⁷ Similarly, the demandingness of a theory is usually a matter of its behaviour overall, rather than any one of its underlying doctrines. Many people seem to have intuitions about demandingness as such, nonetheless.

⁸ I assume that the deontic status of an action is determined by the reasons for or against it. If that is right, insensitivity at the level of reasons implies insensitivity at the level of deontic status. Hence I have formulated the pro-principled intuition in terms of reasons only. Those who disagree with my assumption could amend the discussion of the pro-principled intuition accordingly.

⁹ Of course, there are other possible interpretations of the sentiments expressed when people say 'we should carry on as normal'. One such is that they are concerned with reputation effects, and in particular with not inviting further attacks by rendering the present one effective. However, while such effects are very important, they do not seem to me to capture the full force of the common sentiment that we should carry on as normal when under threat. I think people would believe this even if there were some special feature of the case that prevented the reputation effects from getting any purchase.

¹⁰ Bennett concentrates on problems of formulating the intending/foreseeing distinction, and (especially) the doing/allowing distinction (which he calls the 'making/allowing' distinction). See Bennett [1995: Chapters 4-11]. Kagan emphasises problems in justifying attributing moral significance to either distinction independently of its fit with our intuitions. See Kagan [1989: Chapters 3-4].

¹¹ Elements of this are of course present in Bernard Williams's critique of Utilitarianism. See Williams [1973: 93-100; 108-118]. Some of his remarks on negative responsibility suggest this sort of view, such as his complaint that '... for consequentialism, all causal connexions are on the same level, and it makes no difference, so far as that goes, whether the causation of a given state of affairs lies through another agent, or not' (p. 94). On the other hand, Williams later says of the contrast between action and inaction that 'I doubt whether the sort of dilemma we are considering is going to be resolved by a simple use of this distinction' (p. 109). However, one can view his remarks on integrity as a rival attempt to spell-out the significance of the actor's agency that is allegedly missed by Act Teleology. More recently, Samuel Scheffler has argued that some

version of the doing/allowing distinction is an essential feature of our ethical outlook. See Scheffler [2004: 218-219 and *passim*].

¹² See Woodard [2003]. My views on pattern-based reasons are greatly influenced by Hurley [1989: Chapter 8]. However, she adopts what I'll call in the text below the Willingness Requirement, which I reject.

¹³ Advocates of the idea of pattern-based reasons (not under that name) include: Bacharach [1999]; Hollis [1998: 137-142]; Hurley [1989: Chapter 8; 1991; 1994; 2004]; McClellan [1985]; McMahan [2001]; Regan [1980]; and Sugden [2000]. Critics include: Gruzalski [1986]; Jackson [1997]; Otsuka [1991]; and Tännsjö [1989].

¹⁴ This is an observation about those theories presented as theories of pattern-based reasons: theories of Collective Consequentialism, say, or Plan Consequentialism. However, we might also treat Rule Consequentialism as a theory of pattern-based reasons. Forms of Rule Consequentialism that contain a disaster-prevention rule, such as Hooker's [Hooker 2000: 99], can then be understood as pluralist theories of pattern-based reasons. I shall say a little more about this way of understanding Rule Consequentialism below.

¹⁵ Here I adapt Nozick's formulation of the rationale for symbolic value [1993: 26-35; 41-50].

¹⁶ This constraint comes in different versions, according to the precise specification of willingness required of the other agents. But the core idea stated in the text is shared by these different versions. It is also widely assumed without argument. For example, see: Regan [1980: 124]; Hurley [1989: Chapter 8]; Bacharach [1999]; Sugden [2000]; McMahan [2001].

¹⁷ It is not clear that *everyone driving on the right* is the best pattern *tout court*, though, since *everyone in the world driving on the left* would be just as good, presumably. Like other forms of reasoning, pattern-based reasoning is subject to problems of indeterminacy. I am grateful to Ben Saunders for raising this issue. For the purposes of explaining the worry about recklessness, however, I shall leave it aside.

¹⁸ I assume that we do not need to give a full account of practical relevance to see that some things are practically irrelevant in a given case. In particular, we do not need to decide whether practical relevance is a matter only of the instrumental relation of acts to outcomes, or can also include other ways in which acts can have or be related to value.

¹⁹ In drawing this comparison I have ignored a very important feature of Hooker's view: that it considers the consequences of internalising sets of rules [Hooker 2000: 32]. Rule Consequentialism is most plausibly understood as a theory of pattern-based reasons when the focus is, instead, on the consequences of compliance with rules, since the rules can then be understood as specifying extended patterns of action.

²⁰ The latter point is only a dialectical advantage, since I think that Hooker shows, anyway, how the incoherence objection can be met. See Hooker [2000: 99-102].

²¹ The view of pattern-based reasons sketched here is developed and defended at greater length in Woodard [2007].