

# A New Argument Against Rule Consequentialism

Christopher Woodard

Accepted: 8 August 2007 / Published online: 11 October 2007  
© Springer Science + Business Media B.V. 2007

**Abstract** We best understand Rule Consequentialism as a theory of pattern-based reasons, since it claims that we have reasons to perform some action because of the goodness of the pattern consisting of widespread performance of the same type of action in the same type of circumstances. Plausible forms of Rule Consequentialism are also pluralist, in the sense that, alongside pattern-based reasons, they recognise ordinary act-based reasons, based on the goodness of individual actions. However, Rule Consequentialist theories are distinguished from other pluralist theories of pattern-based reasons by implausible claims about the relative importance of act-based and pattern-based reasons in different cases. Rule Consequentialists should give up these claims. They should either embrace some other pluralist pattern-based view, or reject pattern-based reasons altogether. Note, though, that these arguments apply only to compliance-based, rather than acceptance-based, versions of Rule Consequentialism. This suggests that these two kinds of theory are more different from each other than we might previously have realised.

**Keywords** Act Consequentialism · Rule Consequentialism · Reasons · Rightness · Act-based reasons · Pattern-based reasons · Pluralism · Hooker

Several kinds of objection to Rule Consequentialism are well known. Perhaps the best known is the idea that Rule Consequentialists face a dilemma consisting of choice between practical equivalence with Act Consequentialism on one hand, and incoherence on the other. In my view, Brad Hooker has answered this objection decisively.<sup>1</sup> As a result, discussion of Rule Consequentialism has become more interesting. It tends now to focus on detailed questions about formulation, about the plausibility of its component doctrines, and about the plausibility of its implications.<sup>2</sup>

---

<sup>1</sup>See Hooker (2000, chapter 4).

<sup>2</sup>See, for example, the papers in Hooker et al. (2000).

C. Woodard (✉)

Department of Philosophy, University of Nottingham, Nottingham, Na7 2RD, UK  
e-mail: christopher.woodard@nottingham.ac.uk

I wish to suggest a new sort of criticism of Rule Consequentialism. The criticism will be friendly, since I do not claim that Rule Consequentialism is on the wrong track altogether, but instead that it is wrong in detail. However, I will claim that it is clearly wrong, if only in detail. (I aim to be *friendly but deadly*.) My argument is quite simple, and it depends on two claims. The first is that Rule Consequentialism is best understood as a theory of pattern-based reasons. I shall explain what that means shortly. The second claim is that, thus understood, Rule Consequentialism is distinguished by an account of the importance of such reasons that is very hard to accept. So, believers in Rule Consequentialism should drop this account, and embrace some other theory of pattern-based reasons – or, more radically, reject pattern-based reasons altogether.

I start by considering a formulation of Rule Consequentialism that is some distance from the leading candidate formulations. I'll call this 'Simple Rule Consequentialism'. Starting with it aids exposition, but later I'll consider ways of departing from it. However, there is one important limitation to my argument: it applies only to versions of Rule Consequentialism formulated in terms of compliance with rules, rather than acceptance of them.<sup>3</sup> At the end I'll return to the significance of this point.

## 1 Simple Rule Consequentialism

First, let me explain what I mean by Simple Rule Consequentialism. This is a theory about the rightness of actions, and about moral reasons for or against actions.<sup>4</sup> It claims:

- SRC1 A set of rules is optimal if and only if universal compliance with it would have consequences at least as good as the consequences of universal compliance with any other set.<sup>5</sup>
- SRC2 An action is right if and only if an optimal set of rules permits it.<sup>6</sup>
- SRC3 There is a moral reason for an action if and only if at least one of the optimal rules requires it.
- SRC4 There is a moral reason against an action if and only if at least one of the optimal rules forbids it.

<sup>3</sup>This means that the forms of Rule Consequentialism I shall criticise cannot use Hooker's way of escaping the traditional collapse/incoherence dilemma, since that relies at crucial points on appeals to the costs of acceptance. See Hooker (2000, p. 94). However, I believe they can still escape this dilemma. I shall not try to show this here, since my purpose is not to defend compliance-based versions of Rule Consequentialism (but see note 5 below for some relevant comments). I am grateful to Nicholas Day for discussion of these issues.

<sup>4</sup>To make the claims about reasons more plausible, we have to restrict them to claims about *moral* reasons. Presumably someone could have a (normative) prudential or aesthetic reason for action without that being related in the ways described to a set of rules. So, an interesting feature of Rule Consequentialism is that it seems to be committed to distinguishing in some fairly sharp way between moral reasons and other normative reasons.

<sup>5</sup>Won't Simple Rule Consequentialism contain just one rule: Act Consequentialism? No. We should formulate Rule Consequentialism in a way that captures its appeal, and part of that is to have different implications than Act Consequentialism. Simple Rule Consequentialism may contain rules that are practically equivalent with Act Consequentialism *in a world of universal compliance with those rules* (Regan (1980) aside). But so long as the content of the rules is not a statement of Act Consequentialism, its rules should not be practically equivalent with Act Consequentialism in *other* worlds, including our one. Provided we remember this, we need not adopt acceptance-based Rule Consequentialism to avoid practical equivalence with Act Consequentialism.

<sup>6</sup>For simplicity, I will ignore the possibility of tied optimal sets of rules.

This is almost certainly not the best form of Rule Consequentialism, but it helps to start by considering it. In particular, it takes what many consider to be the wrong tack on two major issues. One of these concerns whether we should evaluate rules in terms of the consequences of *compliance* with them, or instead in terms of the consequences of people *accepting* them.<sup>7</sup> The other concerns the proportion of people whose compliance with, or acceptance of, the rules is relevant to the rules' evaluation. Simple Rule Consequentialism considers compliance by everyone everywhere, and this is widely thought to be the wrong sort of approach.<sup>8</sup> In Section 3 I'll consider some alternative views on this question. At the end I will return briefly to the issue of compliance versus acceptance.

An important feature of Simple Rule Consequentialism is that it includes claims about reasons, as well as claims about rightness. In contrast, most formulations of Rule Consequentialism portray it as making claims only about rightness. This is, I think, a peculiar feature of those discussions; anyway, I shall depart from it. One reason for this is that the possibility of conflict between moral reasons seems worth taking sufficiently seriously for us to want Rule Consequentialism to explain it, should it occur. If reasons do conflict, Rule Consequentialism could explain this only if it makes some claims about what reasons agents have. Note that this rationale does not presuppose that reasons do indeed conflict. It presupposes only that Rule Consequentialism ought to be capable of explaining any such conflict as might occur.

However, even if we set aside the possibility of conflict between reasons, it seems worth formulating Rule Consequentialism such that it includes claims about what reasons agents have. That might not be true if every action that is required, forbidden, or optional had exactly one reason for action associated with it. Were that true, and we knew the deontic status of all the actions we were interested in, it might be said that we would know all the relevant facts. But since it is not true, there is more to explain once we know the deontic status of actions. It remains to be explained what reasons there are for or against each action, even if these reasons do not conflict in respect of any action. So, refraining from formulations of Rule Consequentialism according to which it makes claims about reasons just looks needlessly un-ambitious.

Simple Rule Consequentialism makes claims about reasons and rightness in the following way. It analyses *rightness* in terms of what the whole optimal set of rules permits, and it analyses *reasons* in terms of what individual rules in that set require or forbid. The contrast between reasons and rightness thus depends on the contrast between the implications of one rule and implications of the whole set.<sup>9</sup> On the assumption that there could be conflict between the rules in the set, there could be conflicting reasons, and rightness would then depend on how these conflicting reasons interact.<sup>10</sup>

In sum, Simple Rule Consequentialism has the following features. It evaluates rules in terms of compliance, not acceptance. It takes the simple view that the relevant rate of compliance is 100% of everyone everywhere. Finally, it analyses rightness in terms of the

<sup>7</sup>For discussion of this, see Hooker (2000, pp. 75–80).

<sup>8</sup>For discussion of the considerations pertaining to rates of *acceptance*, see Hooker (2000, pp. 80–85). See also Ridge (2006, *passim*).

<sup>9</sup>This is one way of drawing the contrast between rightness and reasons, and it seems the most natural way for Rule Consequentialism to draw it. Note that Act Consequentialists typically draw it differently, by reference to the contrast between *good in one way* and *good overall*.

<sup>10</sup>One result of this is that there may be no way of judging rightness without relying on intuitions about the upshot of conflicts between reasons. See Hooker (2000, pp. 131–134).

implications of the whole optimal set of rules, and reasons in terms of the implications of individual members of that set.

## 2 Pattern-based Reasons

Having characterised an admittedly simple form of Rule Consequentialism, let me explain why it is best understood as a theory of ‘pattern-based reasons’.<sup>11</sup> Having done that, I will consider whether the same holds true of other forms of Rule Consequentialism, which are more complex and perhaps more appealing.

*Pattern-based reasons* are reasons for or against performing some action, *A*, because of the goodness or rightness of some pattern, *P*, of which *A* is a part. Ordinarily, we tend to think of reasons for action as being *act-based*, which is to say that the reasons for or against an action, *A*, depend on the goodness or rightness of *A* itself. Act-based reasons may be teleological, in the sense that they explain reasons exclusively in terms of goodness. According to Act Consequentialism and Egoistic Decision Theory, all reasons are teleological and act-based. But this restriction to teleology is not part of the idea of act-based reasons itself. For example, to say that some act of torture is intrinsically wrong – and that this provides a reason not to do it – is to assert the existence of a non-teleological act-based reason. The category of act-based reasons is large and inclusive.

Even so, believers in pattern-based reasons claim that act-based reasons are just part of the whole array of reasons.<sup>12</sup> They believe in the existence of *group-based reasons*, in which the pattern in question is performable by a group, or *plan-based reasons*, in which the pattern is performable by an individual agent over time, or both. (Of course, each of these kinds of pattern-based reason could be teleological or non-teleological, too.) Any plausible, full, theory of pattern-based reasons will have to specify which patterns are eligible – that is, which patterns support reasons – and why just these are the eligible ones. However, the bare idea of pattern-based reasons does not specify this. It just uses the idea of a pattern of action in an abstract and schematic way.

It’s pretty clear that we can understand Simple Rule Consequentialism as a theory of teleological pattern-based reasons. Simple Rule Consequentialism claims (SRC3) that there is a moral reason for some action *A* if and only if at least one of the optimal rules requires doing it. Since the optimality of rules is a matter of the value of universal compliance with them, this amounts to explaining any moral reason to do *A* in terms of the goodness of some larger possible pattern of action consisting of universal compliance with the relevant rule. Performance of *A* on the occasion in hand is one part of this pattern. To illustrate this, suppose to begin with that the optimal set of rules includes just one rule. Imagine, for example, that it contains only the rule “always keep promises”. Then Simple Rule Consequentialism claims that there is just one moral reason, which is to keep any promise made; and it claims that this reason exists because of the goodness of the possible pattern of

<sup>11</sup>It is no part of my argument here to defend the existence of pattern-based reasons, though I try to do so elsewhere (Woodard 2007; see also 2002). Instead, I claim here merely that Rule Consequentialism is best understood as a theory according to which there are such reasons. If that is correct, Rule Consequentialists who deny the existence of pattern-based reasons are incoherent, whether or not that denial is correct.

<sup>12</sup>Indeed, one can think of act-based reasons as a species of pattern-based reason, in which the action concerned is identical with the pattern whose rightness or goodness provides the reason. For ease of discussion, however, I shall speak of ‘pattern-based reasons’ in what follows as if the category were properly contrasted with that of act-based reasons.

action consisting of every agent's keeping every promise made. Each agent's reason to keep each one of her promises is that each act of keeping a promise is part of this good pattern. Each such reason is pattern-based, in the sense I've explained.

If we suppose, more realistically, that the optimal set contains several rules, the story is more complicated but the conclusion is the same. Either these rules can conflict with each other or they cannot. If they cannot, then universal compliance with the optimal set of rules amounts to a conjunction of the patterns consisting of universal compliance with each of those rules. This conjunction is itself a pattern, and the reason for complying with each rule would then be that doing so is part of a good pattern – either the conjunctive pattern, or one of the constituent patterns. Again, it would be accurate to portray Simple Rule Consequentialism as claiming that there are pattern-based reasons. What if the optimal set contains rules that could conflict with each other? So long as Simple Rule Consequentialism issues in oughts, and respects the constraint that ought implies can, there must be a way of complying with the whole set of rules, despite the conflict. Perhaps the rules are lexically ordered, or perhaps there are trade-off rules, or perhaps the resolution of conflicts is just left to intuitive judgement in each case.<sup>13</sup> Whatever the details, there will be a possible pattern of action consisting of universal compliance with this set of rules; and Simple Rule Consequentialism will claim that the goodness of this complex pattern provides each agent with reasons to play her parts in it. Once again, we could accurately construe Simple Rule Consequentialism as claiming the existence of pattern-based reasons. Similar remarks go for rules forbidding actions: the pattern, in such cases, is one of universal non-performance of these actions.

This shows that we can think of Simple Rule Consequentialism as a theory of pattern-based reasons. Why think that this is the best way to understand it? Most importantly, because doing so maximises the plausibility of its central ideas. Consider the idea that the reasons for or against an action depend on the consequences of everyone's complying with some rule requiring or forbidding that kind of action. Set aside the practical implications of this idea, and consider it only as a theoretical claim. Compare it, for example, with Act Consequentialism's claim that the reasons for or against an action depend on the causal (and, perhaps, intrinsic) properties of that very action. Leaving aside practical implications, Rule Consequentialism's claim is obscure or mysterious in a way that Act Consequentialism's is not. It is immediately plausible to say that the reasons for or against an action depend on its properties. It is much more obscure, practical implications firmly set aside, to say that the reasons for or against an action depend on the properties of some rule to which it is related. My claim is that we maximise the plausibility of this central component of Simple Rule Consequentialism by understanding it as claiming the existence of pattern-based reasons. No other construal is as charitable.<sup>14</sup>

It is worth pressing Rule Consequentialists for some rationale for their (perhaps implicit) theoretical claims about reasons. Why *should* reasons for action depend on the goodness of rules in this way? It hardly suffices to say that morality is all about rules, since there are plenty of moral views according to which rules play little or no role. If we think of Simple Rule Consequentialism as a theory of pattern-based reasons we can give the following

<sup>13</sup>See note 10 above.

<sup>14</sup>I assume that we should evaluate theories in ethics partly in terms of their practical implications, and partly in terms of the plausibility of their component doctrines. As Parfit writes, "We have intuitive beliefs, not only about which acts are wrong, but also about which principles or theories might be true. So, as well as having plausible implications, a successful principle or theory must be in itself plausible. Only such a principle or theory could *support* our more particular moral beliefs" (2007, p. 264, emphasis in original).

explanation. The concern with the goodness of patterns reflects the idea that actions have *ethically significant parthood properties*. As well as having intrinsic and causal properties, they form parts of larger patterns of action. According to those who believe in pattern-based reasons, these parthood properties can support reasons. If that were so, it would make sense to care about which patterns one is or is not helping to realise when acting. The idea is that this is an intelligible mode of action – just as it is intelligible to care about what one causes, or immediately realises, when acting.<sup>15</sup> In this way, thinking of Simple Rule Consequentialism as a theory of pattern-based reasons provides the best available explanation of why reasons and rightness might depend on the goodness of compliance with rules.

One alternative is to construe Simple Rule Consequentialism as a ‘possibilist’ theory.<sup>16</sup> On this construal, it claims that the reasons for or against an action depend on the *merely possible* consequences of that very action. Though this points us back towards the action itself – which is a plausible place to look for the grounds of reasons – this is not really a good way of understanding Simple Rule Consequentialism. Which of the myriad possible consequences are the ones that ground the reasons? The answer is as follows: the ones that would occur were everyone to comply with the rule in question. This is either just another way of expressing the idea of pattern-based reasons, or it is two mysteries huddled together. First, we need to know why *any* merely possible consequences of actions should matter (remember: independently of the practical implications of this idea). Second, we need to know why *these particular* merely possible consequences should be the ones that matter. The only way of answering these questions, it seems – other than by appeal to the idea of pattern-based reasons – is to appeal to the practical implications of the theory that results. Supposing that these merely possible consequences give rise to reasons might result in a theory that matches our convictions in many cases. But if that is our only ground for believing the theoretical claim about reasons we cannot explain our convictions about cases in terms of the theory.

The idea of pattern-based reasons makes most sense of Simple Rule Consequentialism’s claim that reasons and rightness have to do with the consequences of rules.<sup>17</sup> That is not to say that Simple Rule Consequentialism is convincing when construed this way. After all, it claims that we should comply with rules that would have good consequences if everyone complied with them, even if fewer people will really comply with them. Thus, it claims that we should care about the goodness of *only partly realised* patterns of action, of which our actions are parts. Some people will find that hard to believe. However, there is no better way of making sense of Simple Rule Consequentialism’s concern with the goodness of rules.

<sup>15</sup>Susan Hurley formulated this idea, though in a slightly different form. She writes: “Someone’s reason in acting may ... be that his act bears a constitutive relationship to a valuable form of agency. His contribution to its realization is not a causal one, but that of a part to a whole; it is hardly any less of a contribution, or irrational, on that score. If anything, the relation of part to whole seems more immediate than that of cause to effect” (Hurley 1989, p. 148). Notice that she emphasises being part of a valuable form of agency, whereas I emphasise performing an action that is part of a valuable pattern. One context in which this idea seems particularly salient is when people care about not being complicit in wrongdoing. On complicity, see Kutz (2000).

<sup>16</sup>There is a debate in deontic logic between ‘actualists’ and ‘possibilists’. For a recent discussion, see Zimmerman (1996, chapter 6).

<sup>17</sup>Parfit argues that Kantian Contractualism might justify Rule Consequentialism (2007, pp. 255–260; *passim*). However, despite appearances this is not an alternative rationale for the focus on rules in Rule Consequentialism, since Parfit’s only reason for preferring Rule Consequentialism to Act Consequentialism at the relevant stage in his argument (pp. 255–258) has to do with the two views’ practical implications. It is not a matter of the independent plausibility of the two views’ doctrines.

### 3 More Complex Forms of Rule Consequentialism

As I have described it so far, Simple Rule Consequentialism seems to claim that *all* moral reasons are pattern-based. It claims that all moral reasons have to do with the goodness of compliance with rules, and I have argued that this is best understood as the claim that all moral reasons have to do with the goodness of patterns of action.

What's more, these pattern-based reasons have the following two features. They relate to *very extended* patterns, since Simple Rule Consequentialism evaluates rules on the assumption of compliance by everyone everywhere. Second, these reasons appear to be *insensitive to facts about others' willingness to comply* with the optimal rules. Simple Rule Consequentialism seems to leave no room for any concern with actual levels of compliance, since it seems to claim that what one should do depends only on what would be best if everyone were complying. These two features can make Simple Rule Consequentialism seem objectionably rigid, or un-pragmatic. They seem to make it subject to the charge that it would let the heavens fall rather than adjust its prescriptions in light of information about other agents' recalcitrance. In response to allegations like this, Rule Consequentialists have made several responses. Let us examine three of these.

The first is to point out that the optimal set of rules is likely to contain a *disaster prevention rule*.<sup>18</sup> Without risk of collapse into practical equivalence with Act Consequentialism, Rule Consequentialists can point out that inclusion of a disaster prevention rule in a set of rules is likely to increase the value of that set. So, a set of rules that does not include a disaster prevention rule can usually be improved by addition of such a rule.

How would a disaster prevention rule work? Suppose that the optimal set of rules specifies requisite actions according to kinds of circumstance. We can represent this as follows:

- R<sub>1</sub> In circumstances of kind C<sub>1</sub>, perform (do not perform) an action of the kind A<sub>1</sub>.  
 R<sub>2</sub> In circumstances of kind C<sub>2</sub>, perform (do not perform) an action of the kind A<sub>2</sub>.  
 : :  
 R<sub>n</sub> In circumstances of kind C<sub>n</sub>, perform (do not perform) an action of the kind A<sub>n</sub>.

A disaster prevention rule then says something like this:

- R<sub>d</sub> Where necessary to prevent disaster, break one or more of the other rules.

Of course, a full specification of such a rule would have to spell out what would count as a disaster, and which (if not all) of the other rules could permissibly be broken to prevent one.<sup>19</sup>

Note that endorsing a disaster prevention rule need not constitute a modification to Simple Rule Consequentialism. It is possible that the optimal set of rules, as specified by Simple Rule Consequentialism, would contain such a rule.<sup>20</sup> For this reason we should not expect this point to unsettle the previous conclusion that Simple Rule Consequentialism is best understood as a theory of pattern-based reasons.

<sup>18</sup>For example, see Hooker (2000, pp. 98–99).

<sup>19</sup>Versions of Rule Consequentialism could support exceptionless prohibitions even if they contain disaster prevention rules, so long as they do not claim that *every* other rule could be broken permissibly to prevent disaster.

<sup>20</sup>Simple Rule Consequentialism would include a disaster prevention rule only if its other rules are not practically equivalent to Act Consequentialism in every world. See note 5 above.



On the other hand, it is probably misleading to think of the moral reasons related to the disaster prevention rule as being pattern-based. One could think of every act of preventing disaster (by breaking one of the other rules) as together forming a pattern. One could then think of each such act as forming part of that good pattern, and of this providing a reason to prevent disaster on each occasion. Alternatively, one could think of the ‘disaster prevention rule’ really as consisting of a series of escape clauses from each of the other rules – so that by keeping promises in normal cases, but breaking promises to prevent disasters, say, one would be realising different parts of the *same* complex pattern.<sup>21</sup>

However, these ways of trying to assimilate acts of disaster prevention to other instances of following rules fail to acknowledge the different concerns underlying normal rule following and acts of disaster prevention, for Rule Consequentialists. Rule Consequentialists claim that my reason to keep promises in normal cases has to do with the features of a world of full (or nearly full) compliance with the rule requiring promises to be kept. This is a fairly distant possible world. In contrast, my reason to break a promise to prevent disaster, in exceptional cases, has to do with the features of the *actual* world, with all its faults. It does not have to do with the features of a world of full (or nearly full) compliance even with the rule requiring disasters to be prevented, since some of the occasions on which I need to break a promise (or some other rule) to prevent a disaster will be occasions created by other agents failing to prevent disasters. When disaster is on the cards, it is often because another agent is failing to comply with the optimal rules, and that may include the disaster prevention rule. In attending to these recalcitrant features of the actual world, my root concern is with the harm that my individual action could prevent, given the way the world is. Even if we can describe this as following a rule or participating in a pattern, it really reflects a concern with act-based considerations.

Hence versions of Rule Consequentialism that include disaster prevention rules are *pluralist*, in the sense that they recognise act-based and pattern-based reasons alongside each other. This makes such views more flexible, in the sense that they can respond to the presence of recalcitrant agents, or other nonideal features of the real world. Why should others’ recalcitrance be a problem for Rule Consequentialism? The answer, presumably, has to do with the bad consequences in some circumstances of sticking to rules like  $R_1$  to  $R_n$  in the face of such recalcitrance. If so, pointing out that the optimal set will contain a disaster prevention rule goes a long way to answering the objection.

However, some Rule Consequentialists believe that this does not go far enough, and so they modify the theory in various ways. One such modification is to evaluate rules on the assumption of something other than 100% compliance (or acceptance). Hooker puts the case for this as follows:

A moral code should be suited to the real world, where there is likely to be, at best, only partial social acceptance of, and compliance with, any moral code. An adequate ethic must provide for situations created by people who are malevolent, dishonest, unfair, or simply misguided. In short, for use in the real world, a moral code needs provisions for dealing with non-compliance.<sup>22</sup>

The idea is that the instruction to prevent disaster does not suffice as guidance for dealing with non-compliance. Moreover, since adding extra rules to a code is costly (especially for acceptance-based views), a code that is ideal for 100% compliance or acceptance would not

<sup>21</sup>I am grateful to Alan Carter and others at the BSET Annual Conference for discussion here.

<sup>22</sup>Hooker (2000, p. 80).



contain the rules we need to deal with non-compliance. So it makes sense to evaluate sets of rules assuming something other than 100% compliance or acceptance. Hooker opts for evaluation based on acceptance by 90% of people.<sup>23</sup>

As Hooker acknowledges, the choice of any particular percentage other than 100 as the basis for evaluation of rules will look arbitrary.<sup>24</sup> Largely for that reason, Michael Ridge has made a different proposal. Rather than choosing any one percentage as the basis for evaluation, we should evaluate each set of rules in light of *every* possible rate of compliance (or acceptance) above zero. The best set of rules is then the set that has the highest expected value across these possibilities, treating each possible rate of acceptance as equiprobable. The result, which Ridge calls “Variable Rate Rule Utilitarianism”, is a form of Rule Consequentialism capable of dealing with non-compliance but which does not suffer the sort of arbitrariness associated with picking a single rate below 100%.<sup>25</sup>

Do either of these changes unsettle the claim that Rule Consequentialism is best understood as a theory of pattern-based reasons? Adopting some single rate below 100% as the basis of evaluation certainly does not. It just associates reasons with less extensive patterns, in which not everyone else in like circumstances performs the same kind of action. For example, if we evaluate the rule “if you are well off, contribute 10% of your annual income to aid for the needy” on the hypothesis that 90% of well off people comply with it, we are associating the reason for each such contribution with a pattern in which 90% of eligible contributors do the same.<sup>26</sup> We are still appealing to pattern-based reasons, even though the pattern is somewhat less extensive than if we were to use 100% compliance as the rate for evaluation.

Ridge’s proposal is more difficult to place. The best set of rules on his proposal would not associate moral reasons with any particular pattern of action, but instead with a large set of such patterns, one for each rate of compliance (or acceptance) that is considered. This is, perhaps, a theory of pattern-based reasons in an extended sense. Notice that in respect of rates of compliance nearing zero the patterns concerned may not be any more extensive than the single act under consideration. In this way, Ridge’s proposal seems to marry a concern with the effects of single acts with a concern with the effects of many different successively larger patterns of acts of the same type. Understood that way, it seems to be a way of formulating a pluralist theory of pattern-based reasons, albeit an unusual one.

The final modification I wish to consider is to narrow the scope of Rule Consequentialism. Mulgan distinguishes the “realm of necessity”, which concerns choices affecting the satisfaction of people’s needs, from the “realm of reciprocity”, which concerns choices affecting the satisfaction of people’s goals.<sup>27</sup> In respect of choices affecting needs he advocates Act Consequentialism, while in respect of choices affecting goals he advocates a form of Collective Consequentialism. One could do something similar with Rule Consequentialism in place of Collective Consequentialism. The result would be a

<sup>23</sup>See Hooker (2000, pp. 80–85). See also Ridge (2006, pp. 244–248).

<sup>24</sup>Hooker (2000, pp. 83–85).

<sup>25</sup>Ridge (2006, p. 248). The statement of Variable Rate Rule Utilitarianism in this paper suggests that *zero* compliance should also be included, but Ridge corrected that in his presentation ‘Climb Every Mountain?’, at the University of Reading, 3rd November 2006.

<sup>26</sup>See Hooker (2000, p. 163).

<sup>27</sup>He also identifies the “realm of creation”, which concerns questions about which people to bring into existence. We can set this aside, however, since it does not bear on our present discussion. See Mulgan (2001, pp. 170–172).

theory of pattern-based reasons for the realm of reciprocity, coupled with a theory of act-based reasons for the realm of necessity.<sup>28</sup>

Each of these proposals renders the resulting view better able to respond to recalcitrant agents. Strikingly, the first and third do so in roughly the same way, by combining a concern with the goodness of patterns of action with a concern with the goodness of individual actions. Arguably, Ridge's proposal does this too. These are all ways of being pluralist, in the sense of finding some role for act-based reasons alongside pattern-based reasons – even though these views differ considerably in the ways they do this.

This convergence on pluralism is just as well. *All* believers in pattern-based reasons should be pluralists. As I said earlier, belief in pattern-based reasons makes sense only if actions have ethically significant parthood properties. The basic idea of pattern-based reasons, after all, is that the goodness or rightness of some pattern can provide a reason to perform parts of that pattern. Some find that idea hard to believe. But whether or not this idea is correct, it is very implausible to claim that actions do not have ethically significant properties of other kinds. Most obviously, they certainly have ethically significant *causal* properties. It would be astonishing if causal properties never gave rise to moral reasons. So, everyone should believe that there are some act-based moral reasons. Thus, believers in pattern-based reasons should believe that there are some act-based reasons. Therefore, they should be pluralists.

This provides a second strong reason for understanding Rule Consequentialism as a theory of pattern-based reasons. Understanding it that way explains the appeal of disaster prevention rules, of evaluating sets of rules taking into account non-compliance, of limiting the scope of Rule Consequentialism, and of other similar devices. These are all ways of finding some role for act-based considerations alongside pattern-based considerations, as befits our imperfect world. Since any plausible theory of pattern-based reasons must be pluralist, it is no surprise that Rule Consequentialists have found ways of being so.

#### 4 The Distinguishing Features of Rule Consequentialism

If the argument so far is correct, belief in Rule Consequentialism – at least, belief in plausible compliance-based versions of it – really amounts to belief in a pluralist theory of pattern-based reasons. However, the reverse claim does not hold, since the family of pluralist theories of pattern-based reasons is considerably broader than the family of versions of Rule Consequentialism. So we should ask what distinguishes Rule Consequentialist views from other pluralist theories of pattern-based reasons.

Consider a theory which says that some moral reasons have to do with the goodness of collective actions, and that some have to do with the goodness of individual actions. This is a pluralist theory of pattern-based reasons. It is broadly like the plausible forms of Rule Consequentialism that we have been considering, except that appeal to collective actions replaces appeal to rules. How does that make a difference? It might exclude consideration of patterns consisting of repeated performance of the same sort of action by the same

<sup>28</sup>Of course, Mulgan's own view includes a pattern-based element, since Collective Consequentialism is a pattern-based theory. The point of replacing Collective Consequentialism with Rule Consequentialism is not to introduce pattern-based reasons, but to make his proposal directly relevant to our topic.

agent (unless we were willing to count these as ‘collective actions’ in an extended sense). On the other hand, it would include patterns consisting at least partly of different agents performing different sorts of action. For example, suppose that you and I are walking in some remote place, and we come across an injured person who is barely conscious. It may be that the best that you and I could do in this case would be for one of us to stay with the injured person, trying to prevent him losing consciousness, and for the other to go for help. In other words, the best pattern we could perform might have us doing different things.

For this reason we might say that Rule Consequentialists are distinguished from other believers in pattern-based reasons by virtue of their acceptance of the *Rule Idea*. According to this idea, the only patterns that support pattern-based reasons are ones in which everyone performs the same kind of action in the same kind of circumstance. But why should that be? Why should being part of just that sort of pattern be reason-supporting? On the face of it, the Rule Idea looks implausible, and so we might think this a defect of Rule Consequentialism.

However, this may be to caricature Rule Consequentialism. In particular, it may be to suppose an unduly narrow understanding of the sorts of rules involved. If we allow rules to be somewhat more complex, or less specific, than we have been supposing, then the Rule Idea might look less troublesome. Consider again the example in which we come across an injured person. Though the best pattern we could perform has us doing different things under one set of descriptions, it may be that it could be described in some different way such that we are both following the same rule. For example, we might both be following the rule, “help those in great need, if the cost to you is not prohibitive”. Rules might thus be more encompassing by being less specific. Another possibility is that they might contain complex clauses describing different specific actions. Either way, the Rule Idea might not be as constraining as it seems at first.

For this reason we should probably set aside worries about the Rule Idea. What else distinguishes Rule Consequentialism from other pluralist theories of pattern-based reasons? The feature to which I wish to draw attention is the sort of account Rule Consequentialists have to give of the relative importance of act-based and pattern-based reasons. This account is, I claim, implausibly simple. This is what makes Rule Consequentialist views less credible than other possible pluralist theories.

A preliminary point is that, to be worthy of the name, any version of Rule Consequentialism has to claim that pattern-based reasons are usually more important than act-based reasons (at least, within whatever domain of cases about which Rule Consequentialism is held to be correct). Act-based reasons must not be typically most important, or else the rightness of actions would not normally be a matter of whether the optimal set of rules requires or forbids them. Any act-based reasons that are recognised by a Rule Consequentialist view must work at the margins rather than being typically very important in determining what reasons agents have, or which of their actions are permitted. Some might doubt whether this general predominance of pattern-based reasons is plausible, since it amounts to giving priority to the parthood properties of actions as against their causal (and, perhaps, intrinsic) properties in most cases. However, this might be something on which intuitions just differ, and this sort of doubt is not part of my objection to Rule Consequentialism here.

Instead consider the sort of explanations that Rule Consequentialists can give of why act-based reasons are sometimes more important. The ideas of a disaster prevention rule and

of limiting the scope of Rule Consequentialism each yield a clear explanation of this, so let us consider these explanations in turn.<sup>29</sup>

First consider the role given to act-based reasons by disaster prevention rules. Such rules claim that act-based reasons become important when and because the costs of sticking to the other rules in the optimal set become too great.<sup>30</sup> Now, obviously there is room for disagreement between Rule Consequentialists over how great the costs must be to constitute a disaster. The lower the threshold, the greater is the role given to act-based reasons. But the crucial point is that only one factor is supposed to govern the relative importance of act-based and pattern-based reasons, namely the magnitude of the costs involved in following the pattern-based reasons.

This is implausible for two reasons. The first is that if it were correct, we should be able to produce an ordering of cases according to the costs involved in sticking to the (non-disaster preventing) rules, such that there is no backtracking in the transition from cases in which act-based reasons predominate to cases in which pattern-based reasons predominate, as the costs decrease. However, consider the following two cases. In the first I must choose between telling the truth, in which case a third party will inflict some pain on you, and telling a lie, in which case you will be spared this pain. The second case is just like the first, except that *I* will suffer the pain if I tell the truth, and it will be *greater*. Now suppose that the version of Rule Consequentialism we are considering includes a rule against lying, and that it includes a disaster prevention rule, specified such that the pain you will receive if I lie counts roughly as a disaster. If the cost involved in sticking to the rule is the only relevant factor determining the relative importance of act-based and pattern-based reasons, we must conclude that I should lie in the second case if I should lie in the first. However, I take it that I may be permitted not to lie to protect myself even when I am required to lie to protect you. The fact that the costs of my not lying will be borne by me makes some difference, in addition to whatever difference the size of the costs makes. More generally, our convictions about when it is right to stick to rules and when it is right to make an exception do not seem to reflect concern only with the costs of sticking to the rules.

Second, when we consider the costs of sticking to the (non-disaster preventing) rules we are looking at things from the point of view we occupy when thinking about act-based reasons. Consider again the two cases in which I must decide whether to lie to prevent pain. The cost of sticking to the rule against lying is the pain caused by my not lying. But this is just what grounds my act-based reasons in these cases. This cost is what my action of not lying would cause in the actual world, given the facts about the response of the third party; it is what gives me an act-based reason to lie. However, it is implausible to claim, in effect, that act-based reasons govern the interaction between act-based and pattern-based reasons – *especially* for a Rule Consequentialist, who one would think would prioritise pattern-based

<sup>29</sup>We can understand Ridge's proposal as being pluralist in the relevant sense, as I explained in Section 3. However, if we understand it in that way it does not yield a clear explanation of why act-based reasons should sometimes predominate, since its concern with the consequences of single acts (or very attenuated patterns) is not identifiable with any feature of the final code (in contrast to the devices of a disaster prevention rule, or Mulganesque division of realms). It forms one part of the complex evaluation of each candidate code.

<sup>30</sup>Some versions of Rule Consequentialism might claim that the act-based reasons exist all along, but only become important when a disaster is on the cards; others will claim that they only exist when a disaster is on the cards. The former claim is more plausible, since it is difficult to understand why the causal properties of actions, for example, should support reasons only when a disaster is on the cards. It is somewhat easier to suppose that these properties always support reasons, but that these reasons are important only when a disaster is on the cards.

reasons, if anything. The situation resembles the one we face when there are conflicting reasons associated with different values, such as pleasure and excellence. What governs the interaction of these reasons? If we really do value excellence separately, it would be implausible to answer: “prospective loss of pleasure governs whether pleasure-based or excellence-based reasons are decisive”. But that is analogous to the claim that, when act-based and pattern-based reasons conflict, the costs of acting on the pattern-based reasons determine whether the act-based or pattern-based reasons are decisive. Both of these claims illicitly privilege one of the perspectives from which the conflicting considerations arise.<sup>31</sup>

For these two reasons, it is implausible to think that the device of a disaster prevention rule could capture the relative importance of act-based and pattern-based reasons in an overall economy of reasons. That device points to the costs of following the pattern-based reasons. But other factors seem relevant too, such as who would bear those costs, and whether they would consent to them, and perhaps whether they deserve to bear them. And in any case we should not privilege the perspective from which act-based reasons are salient when thinking about their interaction with pattern-based reasons; at least, not without argument.

Now consider the suggestion derived from Mulgan’s view, that the scope of Rule Consequentialism should be limited, perhaps to choices in which people’s needs are not at stake. This is quite similar to the suggestion we just considered, since it explains the interaction between act-based and pattern-based reasons in terms of the harm caused by following pattern-based reasons. The difference is that the disaster prevention rule expresses the idea that the *amount* of harm is what matters, while the present suggestion is that the *category* of harm is what matters. On the Mulganesque view, if the cost of following pattern-based reasons is harm to the fulfilment of needs, then act-based reasons are more important; but if the cost of following pattern-based reasons is harm to the fulfilment of goals, pattern-based reasons are more important.

Like the disaster prevention rule, this helps Rule Consequentialism to get some important cases right, by allowing it to recognise the importance of act-based reasons in severe emergencies.<sup>32</sup> What’s more, it doesn’t suffer from the defect of illicitly privileging one perspective over the other, since it just assigns each perspective different roles. However we draw the underlying contrast between needs and goals, though, it remains implausible to think that the relative importance of act-based and pattern-based reasons depends on just one factor. To see this, consider another two cases. In the first I can prevent some minor damage to the fulfilment of Smith’s needs only by breaking a rule. In the second I can prevent huge damage to the fulfilment of many people’s goals only by breaking the same rule. Provided the magnitude of the harm in the second case is great enough, it seems that the reasons for breaking the rule must be as strong as in the first case. Once again, it is implausible to claim that a single factor – in this case, the category into which the cost falls – governs the relative importance of act-based and pattern-based reasons.

The general lesson is this. The neatness and simplicity of Rule Consequentialism’s account of the relative importance of act-based and pattern-based reasons distinguishes it

<sup>31</sup>The fact that the disaster-prevention rule embodies a tacit theory according to which the considerations grounding act-based reasons govern the interaction between pattern-based and act-based reasons may be a residue of truth in the old objection that Rule Consequentialism is incoherent.

<sup>32</sup>For example, it gets the right answer in cases in which doing only one’s fair share of disaster prevention would have very bad consequences. In contrast, Murphy’s Collective Principle of Beneficence seems to get this sort of case wrong. See Murphy (2000, pp. 127–134).

from other pluralist theories of pattern-based reasons. I have argued that this makes its account of these matters implausible. If so, Rule Consequentialists should embrace a messier pluralist theory instead, according to which the interaction between act-based and pattern-based reasons depends on several factors and does not fall into any neat pattern. Alternatively, they should give up belief in pattern-based reasons altogether. Either way, they should not remain Rule Consequentialists.

Suppose someone were to develop a pluralist theory incorporating a messier account of the interaction between act-based and pattern-based reasons, and then insist that such a theory still deserved the name 'Rule Consequentialism'. Such a theory would be so different from currently existing forms of Rule Consequentialism that this name might be misleading. However, my objection is not really to the name of the doctrine, but instead to the substance. What we currently understand as Rule Consequentialist views offer only simple accounts of the interaction between act-based and pattern-based reasons. We should reject these accounts.

## 5 Conclusion

The arguments I have given apply specifically to compliance-based versions of Rule Consequentialism, not to acceptance-based versions of the theory. Compliance-based Rule Consequentialism is best understood as a theory of pattern-based reasons; plausible versions of this theory will be pluralist, in the sense that they recognise act-based reasons as well as pattern-based reasons; but these versions are inferior to other pluralist theories of pattern-based reasons that give a messier, more plausible account of the interaction between act-based and pattern-based reasons. So, nobody should believe in compliance-based Rule Consequentialism. They should perhaps believe in some close but less tidy relative of it.

What follows, if anything, about acceptance-based Rule Consequentialism? Though it makes sense to think of compliance with rules as a pattern of action, it is harder to think of acceptance of rules in that way. If we cannot properly think of the goodness of acceptance of rules as the goodness of a pattern, then we cannot properly treat acceptance-based versions of Rule Consequentialism as theories of pattern-based reasons.<sup>33</sup> In one way, that would be a disappointing limitation to my argument. However, it would have an interesting corollary when taken together with my other conclusions. It is worth asking of any Rule Consequentialist theory what its account of reasons is, and why reasons for action should have anything to do with the goodness of rules. If compliance-based versions of Rule Consequentialism are theories of pattern-based reasons, that explains their appeal to the goodness of rules. If, on the other hand, acceptance-based versions of Rule Consequentialism are not theories of pattern-based reasons, then any rationale for their appeal to the goodness of rules must be quite different. If that is correct, these two versions of Rule Consequentialism must be less closely related than we might otherwise have thought. In fact, they would then seem to be quite distinct moral views, with quite different deep foundations.

<sup>33</sup>Brad Hooker suggested to me that there is a sense in which acceptance-based Rule Consequentialism might be pattern-based: it explains reasons in terms of the goodness of *patterns of motivation*. This is a different species of pattern-based theory, which merits further investigation.

**Acknowledgements** For their helpful comments and questions, I am very grateful to audiences at the University of Nottingham and at the Annual Conference of the British Society for Ethical Theory, University of Bristol, 9–11 July 2007. I am especially grateful to Roger Crisp and Nicholas Day.

## References

- Hooker B (2000) *Ideal code, real world*. Clarendon, Oxford
- Hooker B, Mason E, Miller D (eds) (2000) *Morality, rules, and consequences*. Edinburgh University Press, Edinburgh
- Hurley SL (1989) *Natural reasons*. Oxford University Press, New York
- Kutz C (2000) *Complicity*. Cambridge University Press, Cambridge
- Mulgan T (2001) *The demands of consequentialism*. Clarendon, Oxford
- Murphy LB (2000) *Moral demands in nonideal theory*. Oxford University Press, New York
- Parfit D (2007) Climbing the mountain. [http://individual.utoronto.ca/stafforini/parfit/parfit\\_-\\_climbing\\_the\\_mountain.pdf](http://individual.utoronto.ca/stafforini/parfit/parfit_-_climbing_the_mountain.pdf). Cited 21 March 2007
- Regan D (1980) *Utilitarianism and cooperation*. Clarendon, Oxford
- Ridge M (2006) Introducing variable-rate rule-utilitarianism. *Philos Q* 56:242–253
- Woodard C (2002) Group-based reasons for action. *Ethical Theory Moral Pract* 6:215–229
- Woodard C (2007) *Reasons, patterns, and cooperation*. Routledge, New York (in press)
- Zimmerman MJ (1996) *The concept of moral obligation*. Cambridge University Press, Cambridge