UNIVERSITY^{OF} BIRMINGHAM University of Birmingham Research at Birmingham

Act and Rule Consequentialism

Suikkanen, Jussi

DOI: 10.1515/mopp-2023-0075

License: Creative Commons: Attribution (CC BY)

Document Version Publisher's PDF, also known as Version of record

Citation for published version (Harvard): Suikkanen, J 2024, 'Act and Rule Consequentialism: A Synthesis', *Moral Philosophy and Politics*. https://doi.org/10.1515/mopp-2023-0075

Link to publication on Research at Birmingham portal

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

•Users may freely distribute the URL that is used to identify this publication.

Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)

•Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

Jussi Suikkanen* Act and Rule Consequentialism: A Synthesis

https://doi.org/10.1515/mopp-2023-0075

Received August 25, 2023; accepted January 13, 2024; published online February 9, 2024

Abstract: As an indirect ethical theory, rule consequentialism first evaluates moral codes in terms of how good the consequences of their general adoption are and then individual actions in terms of whether or not the optimific code authorises them. There are three well-known and powerful objections to rule consequentialism's indirect structure: the ideal-world objection, the rule-worship objection, and the incoherence objection. These objections are all based on cases in which following the optimific code has suboptimal consequences in the real world. After outlining the traditional objections and the cases used to support them, this paper first constructs a new hybrid version of consequentialism that combines elements of both act and rule consequentialism. It then argues that this novel view has sufficient resources for responding to the previous traditional objections to pure rule consequentialism.

Keywords: act consequentialism; agent neutrality; agent relativity; ideal-world objection; rule consequentialism; rule worship

1 Introduction

Rule consequentialism became a focus of serious philosophical investigation in the 20th century through the works of Brandt (1979), Lyons (1965), Smart (1956), and others.¹ According to a basic version, an action is right if and only if it is authorised by a moral code the general acceptance of which would have the best consequences (hereafter the *optimific code*). Here, the rightness of a particular action is determined indirectly by whether the action is authorised by the optimific code.² This structure makes different versions of rule consequentialism indirect, two-stage ethical theories.

A set of powerful objections to rule consequentialism's indirect structure has also been discovered. Section 2 below first outlines the cases used to illustrate these 9

¹ For more recent notable defences, see Hooker (2000), Mulgan (2006) and Parfit (2011, ch. 16).

² Rule-consequentialist views are thus pattern-dependent as the rightness and wrongness of actions depends on features of cooperative patterns (Perl 2021, ch. 1).

^{*}Corresponding author: Jussi Suikkanen, Department of Philosophy, University of Birmingham, Edgbaston, Birmingham, B15 2TT, UK, E-mail: j.v.suikkanen@bham.ac.uk. https://orcid.org/0000-0002-2804-1280

objections, and Section 3 then explains the objections themselves: the ideal-world objection, the rule-worship objection, and the incoherence objection. The main aim of this paper is to develop a new response to these traditional objections. Section 4.1 first outlines an agent-neutral and an agent-relative theory of value and corresponding consequentialist accounts of practical reasons, and then uses these accounts to respond to the rule-worship and incoherence objections. Section 4.2 relies on the previous two axiologies to formulate a hybrid version of consequentialism that combines elements of both rule and act consequentialism. This new theory gives the optimific codes a new role in the theory of the value of different outcomes, virtue, and agent-relative reasons, rather than in the theory of how rightness relates to good outcomes. Section 4.3 finally shows how the outlined hybrid theory can respond to the ideal-world objection. Section 5 then concludes by responding to an objection to the proposal.

2 Divergence Cases

We first need situations we can call *divergence cases*.³ Imagine you are in a situation in which you can either φ or not- φ and in which the optimific code requires you to φ . Rule consequentialists would claim here that φ ing would be right and not- φ ing wrong. This is a divergence case if and only if not following the optimific moral code would have better consequences than complying with it. Divergence cases are thus situations in which a certain action is required by the optimific code even when doing some other action would bring about a better outcome.

The consequences of an action often depend on whether others are doing similar actions. When we compare the consequences of the general acceptance of moral codes, we must stipulate what level of social acceptance, R%, constitutes general acceptance. When R% of a society has internalized a code, φ ing will have certain consequences just because so many other people are φ ing too. Yet, because a different percentage of people will be φ ing in the actual world, φ ing here can have different consequences, and so we can expect divergence cases to exist. Consider the following divergence case.

Boat: Suppose I and nine other people see ten people fall from the ship's deck into the sea. Each of us still on deck could easily throw one lifejacket to the people overboard. Suppose I throw one lifejacket to the people in the water but then notice that the other people on the deck are doing nothing to help. (Hooker 2000, 164)

³ It should be emphasised that the mere existence of the divergence cases introduced in this section is not a problem for rule consequentialism. In fact, to avoid the collapse into act consequentialism, the rule consequentialists must intend that the view diverges in its implications from act consequentialism (see fn. 8 below). The divergence cases are introduced here merely because they are needed to illustrate the three objections in Section 3.

In Boat, rule consequentialism seems to require you to throw just one lifejacket, because in a world where everyone has internalised the optimific code each person does so.⁴ This is sufficient for saving everyone, and so no further, more demanding obligations with their inculcation costs would be required. However, because in Boat others are not doing their share, following the optimific code has bad consequences and so Boat seems to be a divergence case.

In the opposite kind of cases, an action has bad consequences when everyone acts in that way (and thus it is forbidden by the optimific code), but good consequences in the actual world where others are not doing those actions. Consider:

Fishing. I live in a fishing community that uses traditional small nets to capture fish locally, and everyone captures just enough fish for themselves. If we all begin to use modern ships and large nets, fish stocks will be depleted. But, before others do so, I could use a larger trawler and a bigger net. There would still be enough fish, and I could use the extra income to help the community.

The optimific code here seems to require everyone to use small nets. Yet, before others do so, it would be better for everyone if I used a bigger net. Fishing thus seems to be a divergence case.⁵

Three additional observations can be made about divergence cases. Firstly, many of them are science fiction cases that involve *gremlins* or indestructible devices, *utility landmines*, which create general happiness or misery when a principle is accepted by the overwhelming majority (Podgorski 2018; Rosen 2009). For example, consider:

Gremlin. We live in a world in which a gremlin (i) produces vast amount of happiness if all of us accept a moral rule that requires us to do a jumping jack at noon and (ii) a mild headache for everyone if not all of us accept the rule and yet someone does a jumping jack at noon. As a matter of fact, in this world most people have not internalised the relevant rule.

Rule consequentialism here seems to require you to do a jumping jack at noon in this world even when doing one only brings about mild headaches for everyone.⁶

⁴ Here (and in cases below) I merely want to claim that rule consequentialism *seems* to require you to throw just one lifejacket, and so the case *seems* to be a divergence case. The most plausible versions of rule consequentialism do not have this implication because they include a 'Prevent disasters above all!' principle among the ideal code (Brandt 1992, 151; Hooker 2000, 86, 98–9, and 164). Because of this, not all of the cases discussed here are genuine divergence cases, but at this point these cases are used merely for illustrative purposes and for the sake of simplicity. For the general problems of the disaster-prevention clause, see, e.g. Arneson (2005) and Kahn (2013) and for a discussion of these problems, Woollard (2022). For a discussion of why the disaster-prevention principle cannot help rule consequentialism to avoid the ideal-world objection discussed in Section 3, see Pogdorski (2018, 286). **5** See Perl (2021, 118) and Smith (2010, 419).

⁶ Utility landmines and gremlins are much used in the current literature. Some might argue that such cases are too easy to construct and not particularly convincing. It could furthermore be claimed that these cases include beings who distribute utility based on any arbitrary set of conditions we want

Secondly, not all divergence cases concern the differences between the social acceptance rates in the ideal and real worlds. Consider Judith Jarvis Thomson's case (1976, 206):

Transplant. You, as a surgeon, could save five seriously ill patients by removing the healthy organs of an orphan.

The consequences of the general adoption of a code that authorizes the operation are bad because we would become too anxious to visit a doctor (Parfit 2011, 363), and so the optimific code forbids the operation even if violating that code by operating has better consequences. But, here, nothing depends on how many people are doing such surgeries in the real world as your operating will always have better consequences no matter what others are doing.

Thirdly, similar issues also arise in the real world.⁷ Consider the principles that govern how we should react to (i) aggression, (ii) global poverty, or (iii) climate change. Here too, the same actions can have very different consequences depending on what others are doing. Pacifism has good consequences when others too are pacifists, but bad consequences when they are aggressive (Parfit 2011, 312). Likewise, modest actions of charity or climate change mitigation have excellent consequences when everyone acts in the same way, but they fail to make a difference when others are not doing their share.⁸

to stipulate, which is why it is hard to see what lessons could be taken away from these thought experiments. Furthermore, given how easily gremlins can be adjusted to the details of any theory, it is not even clear whether any view could be gremlin-proof. It could then be suggested that it is enough if our views can deal with real-life examples of our gremlin-free world, which would be enough for these views to be well-adapted for the actual world we inhabit. I thank the anonymous referee of *Moral Philosophy and Politics* for raising this concern with which I have much sympathy. However, at times these thought experiments can be used to clearly illuminate the structural issues of the theories we investigate in the same way as the trolley cases are famously used to illuminate the structural features of various deontic principles. The advantage of these cases is that they enable us to focus on only the morally salient features we want to consider. Furthermore, ideally our theories should be able to capture the essence of right and wrong actions across all worlds rather than merely which actions happen to be right and wrong in our world.

⁷ See Arneson (2005, 239), Brandt (1979, 297; 1983, 42), Card (2007, 244), and Smart (1956, 350–1).

⁸ One way for the rule consequentialists to respond to the traditional objections discussed below would be to try to get rid of the divergence cases. They could, for example, stipulate that an action is right if and only if it is authorised by the optimific code, except in the purported divergence cases in which the right action is to maximize the good. Unfortunately, such responses collapse rule consequentialism into act consequentialism (see, e.g. Brandt 1983, 143; Hooker 2000, 95; and Pogdorski 2018, 281). Another way would be to make room for rules about situations where there is both some nonacceptance and noncompliance (Hooker 2000, 82). If the moral codes were compared with lower levels of acceptance, there would be fewer divergence cases in the actual world and they would only

3 Three Objections

According to the ideal-world objection, the divergence cases, first of all, show that rule-consequentialist theories are extensionally inadequate.⁹ According to rule consequentialism, you ought to comply with the optimific moral code even when doing so (i) merely leads to loss of life (Boat), (ii) requires silly, trivial actions (Gremlin), or (iii) entails that significant goods cannot be obtained (Fishing). The objection then is that these are mistaken moral verdicts. Our intuition in Boat, for example, is that not saving additional lives would be wrong. Thus, the objection is that in many divergence cases the intuitively right action is not to follow the optimific moral code but rather to bring about the best outcome. In addition, the ideal-world objection also offers a diagnosis of where things have gone wrong. The thought is that rule consequentialism must be extensionally inadequate because it focuses on rules that are selected in circumstances that are too different from the circumstances in which they are to be applied.

The advocates of this objection need not, however, claim that you should violate the optimific code in all divergence cases. For example, the optimific promisekeeping principle might well be quite demanding and so require us to keep our promises even when doing so is costly. This is because a demanding promise-keeping principle can foster an atmosphere of trust that is conducive to mutually beneficial cooperation. If this is right, there will be many promise-keeping cases in which following the optimific code will not have the best consequences. Consider an example inspired by Ross (1930, 18):

Promise. You have promised your two children to take them to a park to play. You could take your neighbours' three children instead. This would bring about a little bit more wellbeing than keeping your promise to your own children.

be cases in which intuitively we should follow the optimific code rather than do whatever has the best consequences. This is because, in the actual world, we will be in a context with similar lower levels of social acceptance. Rule consequentialists thus suggest that the moral codes should be compared either at a lower fixed rate of acceptance of, say, 90 % (Hooker 2000, ch. 3.3) or at the optimum rate at which the acceptance of a given code would have the best consequences (Smith 2010). They have also suggested that codes should be compared at all acceptance rates in which case the optimific code either has the best consequences at all rates (Parfit 2011, 319) or on average (Ridge 2006, 248). The hope is that when the codes are compared at lower levels of social acceptance in this way, following the optimific code in the actual world will produce good outcomes too. This would make divergence cases rarer and they would only be situations in which intuitively we should follow the optimific code. For an argument on why this response appears to fail to address the ideal-world objection discussed in Section 3, see Pogdorski (2018, 287–8).

⁹ For different presentations, see, for example, Arneson (2005, 236–7), Brandt (1992, 166), Parfit (2011, 312–20), Perl (2021, ch. 4), and Podgorski (2018).

Intuitively, in Promise you should keep the promise even if this leads to a worse outcome (Hooker 2000, 146; Ross 1930, 18). Similarly, in Transplant, the intuitively right action is to follow the optimific code and not operate even if operating would have better consequences. The advocates of the ideal-world objection can accept these convictions as they need not claim that in all divergence cases violating the optimific principles is right.

The two other objections are the rule-worship and the incoherence objections, which together form a dilemma.¹⁰ Basic rule consequentialism holds that the right action in all divergence cases is to follow the optimific code. Using Derek Parfit's terminology (2011, 38–42) right and wrong furthermore seem to be *reason-implying* properties. Hence, when an action is right there must be sufficiently good reasons to do it, and likewise when an action is wrong there must be sufficiently good reasons not to do that action. Hence, in the divergence cases complying with the optimific code can be right if and only if we have sufficiently good reasons for doing so.

The rule-worship objection is the concern that the rule consequentialists cannot explain why we ought to follow the optimific principles in the divergence cases when doing something else would have better consequences. They can only tell us to follow the optimific principles in these cases, but not give us any good reasons for doing so. The objection then is that this instruction is an objectionable invitation to worship rules for their own sake.

The incoherence objection is the additional concern that if the rule consequentialists try to explain the previous reasons, the resulting accounts will be inconsistent because they require the rule consequentialists to endorse two conflicting theories of reasons. When we consider which moral code should be adopted, the rule consequentialists seem committed to the idea that the reasons for this choice must be based on how good the consequence of the general adoption of the codes are. Yet, if we adopted this account in the divergence cases, we would always have the most reasons not to follow the optimific code, as doing so would bring about more goodness. Rule consequentialists thus need some other account of reasons for the divergence cases, but then that account threatens to conflict with the first account.¹¹

¹⁰ These two objections are not always distinguished from one another. They have been presented in different ways, for example, by Arneson (2005, 236–8), Brandt (1979, 296; 1983, 42; 1992, 150), Card (2007, ch. 3), Darwall (1998, 137–8), Hooker (1995, 28; 2000, ch. 4.3), Kerner (1971, 42–4), Lyons (1965, ch. 4, 1980, 25), Miller (2010), Perl (2021, ch. 3), Railton (1984, 157), Rajczi (2016, ch. 1), Smart (1956, 349–53; 1973, 10), and Wiland (2010, 280).

¹¹ The inconsistency the critics of rule consequentialism have in mind here is not that the theories would imply that agents would sometimes have conflicting reasons. Rather, the claim is that the theories would entail a logical contradiction. The critics are not always clear on what the contradiction would be, but perhaps the idea is that the first theory states that only the total amount of wellbeing can be a reason for an action, whereas the second theory would need to grant that other

4 A New Response

At this point, it is important to observe that the three objections introduced in Section 3 provide two essential desiderata which any defensible version of rule consequentialism would have to be able to meet. Firstly, such a theory must be able to match our moral convictions about the divergence cases so that it can avoid the idealworld objection. This requires being able to explain why in some divergence cases (Boat, Fishing, and Gremlin) we intuitively ought not to follow the optimific code, whereas in others (Transplant and Promise) we intuitively ought to comply with it even if doing so does not have the best consequences. Secondly, the theory must also be able to give a substantial and coherent account of the reasons we have for complying with the optimific code in those cases so that the rule-worship and incoherence objections can be avoided too. Any rule-consequentialist theory that fails to meet these desiderata will be defeated by the previous traditional objections.

This section attempts to outline a new response to those objections which is explicitly aimed at meeting the two desiderata.¹² This response is inspired by Portmore's

12 There are also of course other responses to the three objections. Perl (2021, ch. 2), for example, suggests that we can avoid these objections if we take moral requirements to be demands of the morally best rules where the best rules are (i) most general and (ii) ones that can be accredited with most actual agent-neutral good being caused by actions that the rules classify as morally right. Yet, the problem with this response is that if people only do good actions, this view also counts as best moral rules those rules that permit both good and very bad actions (as the latter actions have no bad consequences in that world) (Chappell 2022). Miller (2010, ch. 5), by contrast, argues that rule consequentialists can avoid the incoherence objection if they (i) take actions to be wrong 'just if the agent would have internalized a rule against it had the people responsible for her moral education ... maximized utility when they engaged in conduct that influenced which rules she internalized', (ii) adopt an act-consequentialist account of practical reasons across the board, and (iii) give up the moral rationalism thesis (see Section 4.2 below). This view is still vulnerable to the ideal world

considerations must be able to provide reasons too. Overall, the rule-worship and incoherence objections boil down to asking for a plausible account of what reason there is to stick to the rules when doing so does not maximise the good (for this point and a different attempt to answer it, see Woodard 2019, ch. 5). In response to these objections, Hooker (1995, 29; 2000, 101) has suggested that rule-consequentialist agents need not have an overarching commitment to maximizing the good, but rather (i) fundamental motivation to do what is impartially justifiable, (ii) a belief that acting according to impartially justifiable rules is impartially defensible, and (iii) a belief that rule consequentialism is on balance the best account of impartially justifiable rules. Whilst this response only addresses the coherence of rule-consequentialist agents and so fails as a response to the rule-worship and incoherence objections (see Card 2007, 246; Miller 2010, 108; and Rajczi 2016, 859–860), Alex Rajczi (2016, 871–2) has reformulated the response as an account of the relevant reasons to stick to the rules when doing so does not maximize the good. Rajczi (2016, 872) also argues that this amended response would still amount to merely asserting that we have sufficient reasons to follow the optimific principles in the divergence cases without explaining what those reasons are.

(2003) dual-ranking act consequentialism that uses two mutually consistent axiologies to formulate a version of act consequentialism that fits our intuitions about agent-centred prerogatives. Section 4.1 similarly outlines two mutually consistent axiologies and uses them to formulate two consistent consequentialist accounts of practical reasons. This section also suggests that the resulting view of reasons can provide a response to the rule-worship and incoherence objections, and so the view meets the second of the previous two desiderata. After this, Section 4.2 formulates two new hybrid versions of consequentialism that combine elements of both act and rule consequentialism, and Section 4.3 finally outlines how these two views can deal with the divergence cases of Section 2. That section thus provides a new response to the ideal-world objection. This means that the proposal can also meet the first desiderata, or so I will argue.

4.1 Two Accounts of Value and Reasons

The first axiology we need for the two new hybrid theories is the traditional *agent-neutral* theory of value, which the rule consequentialists have traditionally used for evaluatively ranking the outcomes of the general adoption of different moral codes.¹³ According to rule utilitarianism, how good an outcome is depends only on the total amount of wellbeing it contains. More recently, rule consequentialists have insisted that the goodness of an outcome may also depend more on how well off the worst-off individuals are and/or the amount of other goods such as fairness, biodiversity, and/or knowledge. Here the selection of the most plausible axiology depends on how well the resulting versions of rule consequentialism can fit our moral convictions and how unified and explanatory the resulting account will be (Hooker 2000, ch. 1 and Sections 2.3–2.5). For our purposes, it does not matter exactly which prima facie plausible agent-neutral axiology we adopt at this point. The axiology we need is thus a previous type of an ordinary rule-consequentialist account of agent-neutral value, call it *value*_{a-n}.

The second axiology we need must be *agent-relative*. Such axiologies are common in the so-called 'consequentialising' literature, where they are used to formulate versions of act consequentialism that can accommodate agent-centred constraints.¹⁴ These versions of act consequentialism agree that we are not always permitted to do whatever makes things go best agent-neutrally, because our own agential involvement

objection type of concerns. For example, perhaps the educators are in circumstances in which helping the poor instead of teaching proper moral principles maximizes utility, in which case unintuitive actions would be right and wrong for the agent. Hence, currently, there is no generally accepted response to the three objections.

¹³ For a discussion of different alternatives, see Hooker (2000, ch. 2).

¹⁴ See Louise (2004), Portmore (2009, ch. 1), and Smith (2003).

(such as our actions of harming or killing others) can sometimes make the resulting outcomes worse relative to us but not relative to others. What is distinctive about such agent-relative axiologies is that the very same aspect of an outcome can be good relative to one agent but either neutral or bad relative to others.

The agent-relative axiology we need has two essential features. Firstly, it needs to accept that whatever qualities (such as general wellbeing) make an outcome agentneutrally good according to the previous axiology (that is, whatever has value_{a-n}), those very same features also make outcomes good relative to every agent. Secondly, it furthermore needs to recognise that there are also other considerations related to the agent, her agential involvement, and especially her rule-following and character traits that can also directly affect how good an outcome is relative to her as an individual.

What are these considerations? To answer this question, we can begin from the idea that sometimes when you act, complying with a specific moral principle can be the end of your action.¹⁵ Hence, the end of my action of handing over a book to my friend can, for example, be to comply with the principle that borrowed books must be returned. Furthermore, as Christine Korsgaard (2005, 22) has argued, the end of an action can be thought to constitute in part the identity of that action – it can make the action the action it is. Giving a book to a friend as a birthday present and giving a book to her to comply with the previous moral principle are, after all, intuitively different actions just because they have different ends.

What I then want to suggest is that it can be good relative to an agent that she has done an action that has as its constitutive end the aim of complying with a moral principle that belongs to the optimific code. That the agent has done such an action would, on this view, be something that can itself make an outcome better relative to her. If what is right were then a function of what is the best relative to you, the previous axiology would often entail that you should follow the optimific code even in the divergence cases.

However, to make the previous simple suggestion compelling, a substantial explanation must be provided of just why it would be good relative to an agent that she has done an action that has as its constitutive end a moral principle that belongs to the optimific code. I believe that such an explanation can be provided if we begin from Brad Hooker's (2002) own rule-consequentialist account of moral virtue instead of from an account of the independent moral value of certain character traits.

According to Hooker (2002), there are two necessary requirements that are jointly sufficient for a given character trait to count as a virtue. By character traits, Hooker (2002, 23) means settled dispositions to act, think, and react in certain ways. These dispositions thus contain behavioural elements, but also motivations, seeing things in a certain light, emotions and reactive attitudes, and ways of reacting to different actions with those attitudes.

¹⁵ Kerner (1971, 49–50) attributes this idea to Urmson (1966).

Hooker's (2002, 38–9) first condition for whether some set of the previous types of dispositions (i.e. a character trait) counts as a virtue is whether the agent having that character trait would be good for everyone. That is, the test is whether having those dispositions that constitute the trait have good consequences for wellbeing generally. Following Hursthouse (1999, 204), Hooker (2002, 34) more precisely emphasises that the trait must have the good consequences in 'a favourable environment', as having virtuous dispositions must be 'standardly' beneficial for 'normal' human beings. Even more importantly, Hooker (2000) also stipulates that 'what makes an environment a favourable one is that the other people in it are virtuous and have strongly negative reactions to those who aren't (34).'

We can then make two observations about this necessary condition (hereafter the *optimality test*). Firstly, Hooker (2002, 28 fn. 14; 2000, 91) understands the adoption of a moral code in terms of developing exactly the previous kinds of dispositions to behave, think, and react in certain ways (that is, in terms of developing a moral conscience with a certain shape). This means that (i) adopting the ideal code in the rule-consequentialist sense and (ii) possessing the character traits that meet the previous optimality test for the status of being a virtue are for him one and the same thing. In other words, the content of the optimific code and the content of the dispositions that meet the first test is identical. This is why Hooker (2002, ch. 3) argues that the most plausible form of virtue ethics in the end collapses into rule consequentialism.

All of this entails crucially that the connection between the optimific code and character traits that count as virtues will not be merely accidental on this view, but rather the optimific code provides virtues the content they must have in order to count as virtuous character traits in the first place. This internal connection makes the optimific code an indispensable element of the hybrid view described below. And, even more importantly, the status of being a virtue is on this view ultimately and fundamentally tied to what is of value according to the previous agent-neutral axiology, namely wellbeing, and so in this respect the proposal described in this section will not lead to introducing a third axiology that would be independent of the other two axiologies introduced in this section.

Yet, according to Hooker, it is not sufficient for being a virtue that a given character trait passes the optimality test. Rather, the trait must also be actually desirable for both the agent herself and also for other people around her. This is because being a 'bad virtue' would be a contradiction in terms (Hooker 2002, 23). The character traits that meet this *desirability test* can be desirable already because of their instrumental value. Trustworthiness, for example, is a trait most happy people tend to have, as the relationships it enables are both an important source of pleasure and constituents of human wellbeing themselves. Furthermore, the traits that meet the optimality test are already desirable for others too because they can feel safe around those who possess them, they can trust those people, and like them more easily too. These traits are thus arguably desirable because they create an atmosphere of trust and mutually beneficial co-operation that leads to general wellbeing (Hooker 2000, Section 4.2).¹⁶ They can also be said to be desirable, for example, because developing them is an achievement which itself is an important personal good (that is, a constituent of an individual's own wellbeing) (Hooker 2002, 24).¹⁷

I will assume that a character trait counts as a virtue if and only if it passes the optimality and the desirability tests. We can then finally formulate the required agent-relative axiology. Thus, according to the resulting agent-relative value theory, (i) any consideration that makes an outcome agent-neutrally good also makes it agent-relatively good and (ii) in addition it is good relative to an agent that she does an action that exemplifies a virtuous character trait that is both optimal in the favourable circumstances and desirable for the agent herself and others in the actual circumstances. Let us call this value *value*_{a-r}.¹⁸

Value_{a-n} and value_{a-r} can then be used to give two corresponding accounts of practical reasons. As consequentialist accounts, these accounts take practical reasons to be determined by the value of outcomes. The suggestion is that we should take the previous two axiologies to provide accounts of the reasons which different subjects have. We first use value_{a-r} to give an account of the practical reasons of ordinary situated agents. On this view, we in the real world have reasons to bring about an outcome to the degree that the outcome contains value_{a-r}, that is, insofar as the action brings about (i) things that make the outcome agent-neutrally good and (ii) doings of actions which exemplify virtuous character traits.¹⁹

¹⁶ It could be also argued that doing an action that has an optimific principle as its end shows respect for the corresponding virtue and that showing such respect is good relative to the agent if respecting the virtue in question in this way is fitting (see Sylvan 2020). It could be added there are *robustly demanding goods*, the respect of which demands having stable character traits that consist of having the relevant optimific principles as ends (see Pettit 2015, ch. 1).

¹⁷ It could furthermore be argued that virtuous traits are also desirable in themselves, but for an objection to this claim see Hooker (2002, 25) and so I will not rely on it below. The reasons listed above for why the optimific character traits are desirable also connect these traits to wellbeing more generally, which is the foundation of the first agent-neutral axiology introduced above. This means that a third independent axiology has not been introduced in this sense either.

¹⁸ $Value_{a-r}$ could also recognise other agent and agent-relative values, but the sparse proposal outlined is sufficient for the current purposes of responding to the three objections.

¹⁹ These reasons are thus not pattern-based in the sense that they would depend on the value of many people acting in some pattern or the agent doing an extended pattern of actions over time, but they are pattern-based in the sense that they depend on the actual desirability of a certain pattern of motivations within the agent (see Woodard 2008, 250 and especially 260, fn. 33). The view can therefore avoid the objections in Dietz (2023, ch. 2). Overall, because the proposal accepts reasons of both (i) and (ii) kind, it is a pluralist view (Woodard 2008, 254).

Value_{a-n} can, by contrast, be used to give an account of the practical reasons of an impartial observer. An impartial observer is a theoretical construct of a person who has no personal ties (and so no partialities) and who can never do anything (and so her own virtuous agential involvement cannot affect the value of outcomes).²⁰ We can then think that how strong reasons an impartial observer has for preferring an outcome is determined by how much value_{a-n} the outcome contains. As a consequence, how much agent-neutral value the outcome of the general adoption of a moral code contains determines how strong the reasons are for the impartial spectator to prefer the general adoption of that code (see Parfit 2011, 372).

We can then make two observations. Firstly, the previous two axiologies are fully consistent with one another. It is true that $value_{a-r}$ recognises one source of value that $value_{a-n}$ does not, but this can be explained, because the function of these two value properties is to provide an account of the reasons which different subjects have. The function of $value_{a-r}$ is to give an account of the reasons of ordinary situated agents, and so this value can be sensitive to how good, relative to these agents, their actions are in exemplifying virtuous character traits. In contrast, because the purpose of $value_{a-n}$ is to give an account of the impartial observer's reasons it simply cannot take such an agent-relative value into account, as the impartial observer never does anything.²¹

Secondly, because these two axiologies and the reasons they ground are consistent with one another, the hybrid view I will formulate on the basis of them will not be vulnerable to the incoherence objection. Furthermore, value_{a-r} also provides a substantial, virtue-based account of what can make it good, relative to an agent, that she does an action that has an optimific principle as its end. This account is based on the two tests, closely connected to wellbeing, for why certain character traits such as honesty, beneficence, and so on count as virtues. For this reason, the view introduced below will never ask us to comply with the optimific moral code just for the sake of it, and so the view can avoid the rule-worship objection too.

4.2 A New Hybrid Version of Consequentialism

We now have the resources required for a new hybrid version of consequentialism that combines elements of both act and rule consequentialism. The key innovation of

²⁰ For discussions, see, for example, Harman (2000) and Mill ([1871] 1998, ch. 2, para. 18).

²¹ The proposal thus avoids the logical inconsistency that threatens traditional rule-consequentialist accounts (see fn. 11 above). It claims consistently that, for all agents, only the total amount of wellbeing and that actions are done from virtuous character traits make outcomes good relative to the agent and thus be reason-providing. It is just that, by stipulation, the impartial observer can never actually be in a position to have reasons of the second kind.

this view is that it gives the optimific code a different role. In this theory, that code is not used to explain how rightness relates to good outcomes but rather to explain value, virtue, and the agent-relative reasons associated with virtuous actions.

We can call the new view hybrid consequentialism. This view rejects the idea that right and wrong are a function of what is authorised by the optimific principles, and so in this sense it is not a rule-consequentialist view. Yet, the first positive element of hybrid consequentialism is an account of practical reasons, which combines elements of both act and rule consequentialism. It is act consequentialist in the sense that what an agent has most reason to do (and what she ought to do) is determined directly by which action brings about most (agent-relative) value (i.e. value_{a-r}). It is, however, also rule consequentialist in the sense that the optimific rules can affect how good the outcome of a given action is relative to the agent. This is because, as explained, that an action with a certain optimific principle as its end has been done can be good relative to the agent insofar as doing the action thereby exhibits a virtuous character trait that passes both the optimality and desirability tests. On this view therefore, the optimific code can affect how good the consequences of different actions are relative to the agent, and so also in part determine what the agent has most reason to do and what she ought to do. A ruleconsequentialist element is thus also built into this view's act-consequentialist theories of value, reasons, and ought. On this view thus both what reasons an agent has and what she ought to do is determined by how much value_{a-r} her actions bring about.

The second element of hybrid consequentialism then is that, on this view, which actions are right and wrong is also determined by what maximizes value_{a-r}. Thus, according to hybrid consequentialism, an action is right if and only if it brings about at least as much value_{a-r} as any other action available for the agent. Yet, even if this view is actconsequentialist in structure in this way too, it is not a version of pure act consequentialism, because according to it an agent's reasons (what she ought to do) and which actions are right and wrong all still partly depend on the optimific principles. This is because on this view that an agent does an action that has an optimific principle as its end can itself make the outcome of that action good relative to the agent, and so such principles can affect reasons, what one ought to do, and what is right and wrong.²² Here the ruleconsequentialist element of the view is again built into the view's theories of value, reasons, and ought, and thereby even into the view's account of right and wrong actions.

Let me conclude this section by making one observation about hybrid consequentialism. One initially attractive feature of the view is that given how the view is

²² This feature enables the view to respond to the classic objection to act consequentialism according to which act-consequentialist views cannot take into account the importance of stable commitments and dispositions to act in moral ways (see Railton 1984). It also enables the view to recognise the inculcation costs and expectation effects of principles that are often emphasised by the rule consequentialists (Hooker 2000, 78–80).

formulated, hybrid consequentialism entails the strong *moral rationalism* thesis. Strong moral rationalism is, roughly, the view that we always have decisive reasons to do what is right and to avoid actions that are wrong.²³ Hybrid consequentialism endorses this claim because, according to it, an action that is right must always be an action that you have most reason to do. This is because both rightness and what you have most reason to do is, on this view, determined by what brings about most value_{a-r}.

We have then a new type of a hybrid-consequentialist theory on the table. The final question we thus need to consider is whether it can provide a response to the ideal-world objection.

4.3 The Hybrid Theory, the Ideal-World Objection, and the Divergence Cases

This section returns to the divergence cases of Section 2 by investigating whether hybrid consequentialism can match our convictions about these cases.

There seems to be two kinds of divergence cases. What is distinctive about Boat, Fishing, and Gremlin is that we intuitively should not comply with the optimific code in them. In Boat we should throw additional lifejackets, in Fishing we should use a bigger net, and in Gremlin not do a jumping jack at noon even if these actions are forbidden by the optimific code. By contrast, even if Transplant and Promise are divergence cases, we intuitively ought to follow the optimific code in them. Thus, in Transplant we ought to not operate and in Promise we should take our own children to the park even if there are actions that would have agent-neutrally better consequences.

Let us consider the first kind of cases first, starting from Boat. To see what implications the hybrid view has in Boat, we first need to compare how much value_{a-r} the outcomes of different actions would contain in it. In this case, you can either throw just one lifejacket or continue throwing more lifejackets to save more lives. If you throw just one lifejacket, the consequences are that (i) one life is saved, (ii) nine people die, and (iii) you have done an action the end of which is to comply with a life-saving principle that belongs to the optimific code. It can then be argued that doing an action that has this principle as its end is good relative to the agent because it exhibits the trait of doing your fair, equal share, which perhaps could be suggested to be a virtuous trait. This is because this trait seems to pass both (i) the optimality test and also (ii) the desirability test (because having the trait of doing your fair share is, for example, conducive to an atmosphere of smooth, predictable co-operation). Yet, the consequence of throwing additional lifejackets is that all 10 people are saved, and so arguably here throwing

²³ For a recent defence of moral rationalism, see Portmore (2011, ch. 2) and Van Roojen (2010). The thesis is not, however, universally accepted (see, for example, Williams 1995).

additional lifejackets brings about more value_{a-r} than following the optimific code even when we consider the value of acting from virtuous character traits.

This means that in Boat hybrid consequentialism entails that you have most reason to keep throwing additional lifejackets, that you ought to do so, and that this action would also be the right thing to do. Hybrid consequentialism can thus match our moral convictions about this case insofar as throwing additional lifejackets intuitively is the right thing to do.

Hybrid consequentialism has similar consequences in Fishing and Gremlin. In Fishing, the view can claim that even if using small nets is good relative to the agent because it has an optimific principle as its end and so the action exhibits a virtuous character trait (of, say, caring about the environment), the outcome of using larger nets still has more value_{a-r} given all the agent-neutrally good consequences that follow from doing so. It follows that using larger nets is, according to hybrid consequentialism, what you have most reason to do, what you ought to do, and the right thing to do, which again is the intuitively right verdict.

In Gremlin, having the relevant optimific moral principle that requires doing a jumping jack at noon as an end of your actions cannot, however, be claimed to constitute a disposition that would count as a virtuous character trait. It can, of course, be argued that having this character trait promotes general wellbeing in the favourable circumstances in which others share it, and so the trait passes the optimality test. However, having this trait is not desirable for the agent or for others in the actual circumstances in which others do not share it (it merely produces headaches for everyone) and so the trait fails the desirability test. This means that doing a jumping jack at noon with the optimific jumping-jack principle as its end cannot make the outcome of that action any better relative to the agent in the real world, given that the agent would not be acting from a virtuous character trait. As a consequence, hybrid consequentialism entails that you have no reason to do a jumping jack at noon when others have not internalised the jumping-jack principle, and so it is not what you ought to do or the right thing to do.²⁴

²⁴ It could be argued that it is also possible to construct gremlin cases against hybrid consequentialism. For example, we could imagine a world in which a Gremlin* (i) produces vast amounts of happiness whenever we act based solely upon agent-neutral reasons, and (ii) produces mild headaches for everyone whenever anyone acts on agent-relative reasons. I thank the anonymous referees of *Moral Philosophy and Politics* again for raising this concern. In this scenario, it is true that we would much more often have most reason not to follow the optimific principles in the divergence cases as the additional happiness produced by Gremlin* would more often outweigh the agentrelative goodness produced by acting from the virtuous character traits. But, to me, that sounds like the right conclusion to draw in such a scenario, and so it's not clear whether this scenario could be used to make the ideal-world objection to hybrid consequentialism. Also, given the flexibility of hybrid consequentialism (see Section 5 below) it will be more difficult to construct a version of the gremlin case that would be effective against the view.

What about the divergence cases in which we intuitively ought to follow the optimific principles? In Transplant, it can be argued that when you refuse to operate, you have as an end of your action an optimific 'Do not kill!' principle. It can also be argued that having such a principle as an end constitutes a virtuous character trait – the trait of treating each person as an end in themselves, which could be said to fall under the personal virtue of justice.²⁵ This trait passes the optimality test because if most people lacked it in the society we would be too anxious to see a doctor, which would have catastrophic consequences (Parfit 2011, 363). Having this trait also passes the desirability test because in the real world it enables us to feel safe in knowing that we will not be sacrificed for the sake of others. This trait thus seems to pass Hooker's two tests for virtuous character traits.

To reach the intuitive verdict that we ought not to operate in Transplant, hybrid consequentialism would then have to insist that not killing the orphan, merely letting the five individuals die, and doing an action that exhibits the virtuous character trait is better relative to the agent than killing the orphan and saving the five. With this axiology, the defenders of hybrid consequentialism can argue that not operating maximises value_{a-r} and so it is what you have most reason to do, what you ought to do, and the right thing to do.

Finally in Promise, hybrid consequentialism can recognise that acting from the optimific promise-keeping principle exhibits the virtuous character trait of trust-worthiness and so the outcome of keeping your promise and taking your own children to the park is better relative to you than taking your neighbour's children. Hybrid consequentialists can thus claim that keeping the promise is what you have most reason to do, what you ought to do, and also the right thing to do.²⁶

Overall, this means that hybrid consequentialism seems able to match our moral convictions about the divergence cases we examined. The point of these cases was originally to illustrate the ideal-world objection by showing that rule consequentialism is extensionally inadequate – that due to structural reasons rule consequentialism has implausible ethical consequences. Given that hybrid consequentialism matches our convictions in these cases, the view thus seems able to avoid the ideal-world objection too.

5 Conclusion, an Objection, and a Response

The ideal-world objection, the rule-worship objection, and the incoherence objection are powerful objections to rule consequentialism's indirect structure and have often been

²⁵ See Kamm (1992).

²⁶ That hybrid consequentialism entails that you ought to follow the optimific principles in these last two divergence cases is significant, because it shows that the view does not collapse into pure act consequentialism.

considered to be decisive. This paper's aim has been to outline a new rule-consequentialist response to these objections. The basic approach has been to highlight that to avoid the objections, the rule consequentialists should accept the following claims:

- 1. The optimific code is the one that an impartial spectator would have most reason to prefer. The spectator's axiology and the reasons it grounds are agent-neutral.
- 2. As ordinary situated agents, our own reasons are based on the agent-relative value of outcomes. Everything that has agent-neutral value has agent-relative value as well, but so does compliance with the optimific code, at least in certain cases.
- 3. Compliance with the optimific code has agent-relative value when the end of the action is to comply with a principle of that code in a way that constitutes exhibiting a virtuous character trait. For a trait to count as virtuous, it must have good consequences for general wellbeing in the favourable circumstances (the optimality test) and having the trait must also be desirable for the agent and for others in the actual circumstances (the desirability test).

I have then suggested that hybrid consequentialism with this structure can avoid the three traditional objections to pure rule consequentialism. Such a theory can both provide substantial accounts of our reasons to comply with the optimific code in the divergences cases and match our moral convictions about those cases too. This is possible because hybrid consequentialism does not rely on the optimific code in explaining how rightness relates to good outcomes, but rather it uses the optimific code in the theory of virtue, value, and agent-relative reasons.

I finally want to conclude by addressing an important concern one might have at this point. It is true that the previous proposal is built on several different elements – an agent-neutral axiology, an agent-relative axiology, and a rule-consequentialist account of virtue – and so there is a degree of structural complexity. I have also not fully specified all these elements, which leaves room for flexibility. For example, if doing actions that exhibit virtue has little agent-relative value, this view entails that we should almost never comply with the optimific codes in the divergence cases, which would collapse the hybrid view into pure act consequentialism. By contrast, if we assume that acting from virtuous character traits has a lot of agent-relative value, the hybrid view will entail that we should comply with the optimific code in almost all divergence cases, which would collapse the view into pure rule consequentialism. And, in between these extremes there is a whole spectrum of different views of the agent-relative value of different virtuous actions that would have different consequences for the divergence cases.

This flexibility is arguably problematic for two reasons. Firstly, Hooker's original ambition was to show that rule consequentialism 'does a better job than its rivals of matching and tying together our moral convictions' (Hooker 2000, 104). It could be argued that because of the previous flexibility, hybrid consequentialism fares worse in this respect than, for example, Hooker's own version of pure rule consequentialism, which could be claimed to be more fully specified and hence better able to tie our moral convictions together. Secondly, it could also be argued that even if these elements help respond to the objections, it is not clear that they fit together in a way that would result in a unified and compelling way of thinking about morality generally.

I have three responses to these important concerns. Firstly, many of the traditional pure forms of rule consequentialism are not fully specified in all respects either. Hooker (2000, 43–5) himself, for example, builds his version of rule consequentialism around 'some modest form of an objective-list account of wellbeing' without specifying what goods belong to the list or what weights they have, and he also gives some priority to the wellbeing of the worst off without specifying exactly how much. For this reason, his account too has similar flexibility as the hybrid views outlined above. This, however, is not a problem because we can then use the reflective equilibrium method, together with our moral convictions about cases, to further specify the constituents of wellbeing, the amount of priority we should give to the worse off, and also in this case how good relative to an agent it is that she acts from virtuous character traits (see Hooker [2000, 15]). If this is right, the version of hybrid consequentialism outlined above should be able to tie together our moral convictions at least as well as Hooker's own pure form of rule consequentialism.

Secondly, there is at least one important respect in which the outlined hybrid view is better able to match and tie together our moral convictions than the existing pure forms of rule consequentialism. As we saw in footnotes 4 and 8, even if Hooker's own amendments to traditional rule consequentialism (comparing codes at lower levels of acceptance and including an 'Avoid disasters above all!' principle in the ideal code) enable the view to match our moral convictions about some divergence cases, we can still construct new cases in which Hooker's view still has objectionable consequences. This means that the traditional pure forms of rule consequentialism simply cannot match and tie together all our moral convictions, and so they are still vulnerable to the ideal-world objection. One advantage of the new hybrid view over the more traditional pure forms of rule consequentialism then is that in addition to matching our moral convictions elsewhere just as well as those theories, this new view can also match and tie together our moral convictions concerning all the divergence cases too.

Finally, even if it is true that the view is structurally more complex and involves additional elements, these elements do seem to fit together in a way that can be developed into a unified, compelling moral view. Hybrid consequentialism, after all, begins from the fundamental agent-neutral value of general wellbeing and the ideal of making things go best together (Hooker 2000, 1). Yet, it also recognises that in addition to the previous type of general wellbeing acting from virtuous character traits can also make an outcome good relative to us too. Hybrid consequentialism furthermore holds that in order to count as a virtuous character trait, having the trait must amount to the internalization of an element of a moral sensitivity which both (i) produces wellbeing in the favourable circumstances in which others too share the trait and (ii) is also desirable in the real world as well, due to its tendency to promote the wellbeing of the agent and others around her here too. This suggests that the different elements of the proposed hybrid view can be linked together through the connection they have to general wellbeing and the character traits that promote it, and so the hybrid view seems to offer a new, genuinely unified way of thinking about morality that deserves to be explored further.

Acknowledgements: I thank the Choice Group at the London School of Economics, Aaron Salomon, and the anonymous referees of *Moral Philosophy and Politics* for helpful comments on the earlier drafts of this manuscript.

References

- Arneson, R. 2005. "Sophisticated Rule Consequentialism: Some Simple Objections." *Philosophical Issues* 15: 235–51.
- Brandt, R. 1979. A Theory of the Good and the Right. Oxford: Oxford University Press.
- Brandt, R. 1983. "The Real & Alleged Problems of Utilitarianism." Hastings Center Report 13 (2): 37-43.
- Brandt, R. 1992. Morality, Utilitarianism, Rights. Cambridge: Cambridge University Press.
- Card, R. 2007. "Inconsistency and the Theoretical Commitments of Hooker's Rule Consequentialism." Utilitas 19 (2): 243–58.

Chappell, R. 2022. Objections to Rule Consequentialism. *Philosophy, et Cetera*. https://www.philosophyetc. net/2022/02/objections-to-ruleconsequentialism.html (accessed January 10, 2024).

- Darwall, S. 1998. Philosophical Ethics. Boulder: Westview Press.
- Dietz, A. 2023. "Pattern-Based Reasons and Disasters." Utilitas 35 (3): 131-47.
- Harman, G. 2000. "Moral Agent and Impartial Spectator." In *Explaining Value and Other Essays in Moral Philosophy*, 181–95. Oxford: Oxford University Press.
- Hooker, B. 1995. "Rule Consequentialism, Incoherence, Fairness." *Proceedings of the Aristotelian Society* 95 (1): 19–36.
- Hooker, B. 2000. Ideal Code, Real World. Oxford: Oxford University Press.
- Hooker, B. 2002. "The Collapse of Virtue Ethics." Utilitas 14 (1): 22-40.
- Hursthouse, R. 1999. On Virtue Ethics. Oxford: Oxford University Press.
- Kahn, L. 2013. "Rule Consequentialism and Disasters." Philosophical Studies 162 (2): 219–36.
- Kamm, F. M. 1992. "Non-Consequentialism, the Person as an End-In-Itself, and the Significance of Status." *Philosophy and Public Affairs* 21 (4): 354–89.
- Kerner, S. 1971. "The Immorality of Utilitarianism and the Escapism of Rule-Utilitarianism." The Philosophical Quarterly 21 (82): 36–50.
- Korsgaard, C. 2005. "Acting for a Reason." Danish Yearbook of Philosophy 40 (1): 11-35.
- Louise, J. 2004. "Relativity of Value and the Consequentialist Umbrella." *The Philosophical Quarterly* 54 (217): 518–36.
- Lyons, D. 1965. Forms and Limits of Utilitarianism. Oxford: Oxford University Press.

Lyons, D. 1980. "Utility as a Possible Ground of Rights." Noûs 14 (1): 17-28.

Mill, J. S. [1871] 1998 In Utilitarianism, edited by R. Crisp. Oxford: Oxford University Press.

- Miller, D. E. 2010. "Mill, Rule Utilitarianism, and the Incoherence Objection." In *John Stuart Mill and the Art of Life*, edited by B. Eggleston, 94–116. Oxford: Oxford University Press.
- Mulgan, T. 2006. Future People. Oxford: Oxford University Press.
- Parfit, D. 2011. On What Matters, Vol. 1. Oxford: Oxford University Press.
- Perl, C. 2021. "Solving the Ideal World Problem." Ethics 132 (1): 89–126.
- Pettit, P. 2015. The Robust Demands of the Good. Oxford: Oxford University Press.
- Podgorski, A. 2018. "Wouldn't it Be Nice? Moral Rules and Distant Worlds." Noûs 52 (2): 279-94.
- Portmore, D. 2003. "Position-Relative Consequentialism, Agent-Centred Options, and Supererogation." *Ethics* 113 (2): 303–32.
- Portmore, D. 2009. "Consequentializing." Philosophy Compass 4 (2): 329-47.
- Portmore, D. 2011. Commonsense Consequentialism. Oxford: Oxford University Press.
- Railton, P. 1984. "Alienation, Consequentialism and the Demands of Morality." *Philosophy and Public Affairs* 13 (2): 134–71.
- Rajczi, A. 2016. "On the Incoherence Objection to Rule-Utilitarianism." *Ethical Theory and Moral Practice* 19 (4): 857–76.
- Ridge, M. 2006. "Introducing Variable Rate Rule Consequentialism." *The Philosophical Quarterly* 56 (223): 242–53.
- Rosen, G. 2009. "Might Kantian Contractualism Be the Supreme Principle of Morality?" *Ratio* 22 (1): 78–97. Ross, W. D. 1930. *The Right and the Good*. Oxford: Clarendon Press.
- Smart, J. J. C. 1956. "Extreme and Restricted Utilitarianism." The Philosophical Quarterly 6 (25): 344-54.
- Smart, J. J. C. 1973. "An Outline of a System of Utilitarian Ethics." In Utilitarianism: For & Against, edited by
- J. C. C. Smart, and B. Williams, 1–76. Cambridge: Cambridge University Press.
- Smith, M. 2003. "Neutral and Relative Value after Moore." Ethics 113 (3): 576–98.
- Smith, H. 2010. "Measuring the Consequences of Rules." Utilitas 22 (4): 413-33.
- Sylvan, K. 2020. "Respect and the Reality of Apparent Reasons." *Philosophical Studies* 178 (10): 3129–56.
- Thompson, J. J. 1976. "Killing, Letting Die, and the Trolley Problem." *The Monist* 59 (2): 204–17.
- Urmson, J. O. 1966. "Utilitarianism." Arnold Isenberg Memorial Lecture delivered at Michigan State University, January 21.
- Van Roojen, M. 2010. "Moral Rationalism and Rational Amoralism." Ethics 120 (3): 495-525.
- Wiland, E. 2010. "The Incoherence Objection in Moral Theory." Acta Analytica 25 (3): 279-84.
- Williams, B. 1995. "Internal Reasons and the Obscurity of Blame." In *Making Sense of Humanity, and Other Philosophical Essays*, 35–45. Cambridge: Cambridge University Press.
- Woodard, C. 2008. "A New Argument against Rule Consequentialism." *Ethical Theory and Moral Practice* 11 (3): 247–61.
- Woodard, C. 2019. Taking Utilitarianism Seriously. Oxford: Oxford University Press.
- Woollard, F. 2022. "Hooker's Rule Consequentialism, Disasters, Demandingness, and Arbitrary Distinctions." *Ratio* 35 (4): 289–300.