

ARTICLES

Does the Patterned View Avoid the Ideal Worlds Objection?

Benedict Rumbold 

Department of Philosophy, University of Nottingham, Nottingham, UK
Email: Benedict.Rumbold1@nottingham.ac.uk

Abstract

Can we formulate a moral theory that captures the moral significance of patterns of group behaviour we cannot affect through our own action while at the same time avoiding the so-called ‘Ideal Worlds’ objection? In a recent article, Caleb Perl has argued that we can. Specifically, Perl claims that one view that does so is his Patterned View: roughly, you ought to act only in accordance with that set of sufficiently general rules that has optimal moral value (Perl 2021: 98). The Patterned View undoubtedly constitutes a welcome contribution to our existing set of moral theories. However, does it avoid the Ideal Worlds objection? In this article, I argue ‘no’.

Keywords: Patterned View; Ideal Worlds objection; Rule Consequentialism; Normative Ethics

In a recent article, Perl (2021) presents a new version of rule consequentialism. In a similar vein to positions endorsed by writers like Richard Miller (2009), Perl’s theory looks to be both resolutely pattern-dependent (it grounds moral facts in facts about patterns of group behaviour) yet also hypothetical-free (it only considers *actual* patterns).¹ According to his

Patterned View: you ought to act only in accordance with that set of sufficiently general rules that has optimal moral value. (Perl 2021: 98)

Perl champions the Patterned View on several grounds. Most strikingly, however, he claims that the Patterned View offers a ‘near-Pareto improvement’ over traditional forms of rule consequentialism (i.e. those which engage in hypotheticals) – as well as related pattern-dependent views – in that, unlike those views, it avoids the Ideal Worlds objection (Perl 2021: 117; cf. p. 90; p. 95).

In this article, I contest this latter claim. The Patterned View may have many merits. However, as far as I can see, it does not avoid the Ideal Worlds objection. In that specific

¹Note: here Perl prefers the term ‘counterfactual’. However, I find this term can create confusion. Often ‘counterfactual’ can seem to imply conditionals with contrary-to-fact antecedents. But the world considered by traditional rule consequentialists need not be contrary to the facts. Sometimes the world they consider might be identical to the facts. So I prefer ‘hypothetical’.

regard, therefore, we cannot say it offers a near-Pareto improvement over those pattern-dependent views that engage in hypotheticals.

Clarification on this point matters for at least two reasons. First, it helps us gain a firmer grip on the (relative) merits of the Patterned View. Perhaps more importantly, however, it also helps us to gain a better understanding of the scope and mechanics of the Ideal Worlds objection. It may be that proponents of other pattern-dependent views – for example, contractualists – believe they can avoid the Ideal Worlds objection by eliminating hypotheticals. This, I argue, is not true. Rather, in line with what I take to be Podgorski's view (2018), what matters is whether such theories ground moral facts in facts about patterns of group behaviour too distant to be activated by any of our actions (or too distant for our actions to place a burden on others to act in certain ways with regards to them).

The article progresses as follows. In Section One, I give a brief overview of the Patterned View. In Section Two, I set out the Ideal Worlds objection. In Section Three, I argue that, on a standard reading of both Perl's theory and the Ideal Worlds objection, the Patterned View looks vulnerable to the objection. In Section Four, I consider and reject various counters a proponent of the Patterned View might offer in response. I conclude in Section Five.

I. The Patterned View

According to the

Patterned View: you ought to act only in accordance with that set of sufficiently general rules that has optimal moral value. (Perl 2021: 98)

One way we can understand the Patterned View is in contrast to what Perl takes to be its primary rival: traditional rule consequentialism. Like traditional rule consequentialism, the Patterned View is a *principled* view: it sees the demands of morality as the demands of those rules that are morally best. Where it differs from such views is in its analysis of what constitutes the morally best rules.

According to at least one traditional rule consequentialist position, endorsed by Hooker, the morally best rules are those rules that, if generally internalized, would have as great an expected overall value as any alternative set of rules (Hooker 2002: 32). According to the Patterned View, the morally best rules meet the following requirements:

Generality Requirement: the morally best rules are sufficiently general.

Consilience Requirement: the moral value of a rule R is everything actual that is agent-neutrally good or bad to the extent it is caused by actions that R classifies as morally right.

The Consilience Requirement is perhaps where the contrast with traditional rule consequentialism is felt most acutely. For example, on Hooker's rule consequentialism, if we want to understand which rules are optimific we first need to consider a hypothetical: namely, a world in which R is generally internalized by the overwhelming majority of everyone, everywhere, in each new generation. We then need to ask: what would the expected overall value of R be in that world? If it would be equal to or greater than some competitor rule, it is morally best. If not, the competitor is best.

According to the Consilience Requirement, no hypothetical is necessary. If we are to ascertain whether R is optimific, we need only look at the overall value (good minus

bad) caused by actions that R classifies as morally right. If the overall value of those acts is equal to, or greater than, that generated by actions that some competitor rule classifies as morally right, it is morally best. If not, the competitor rule is best.

For example, let us say – borrowing a case from Perl – that we are trying to decide which of the following two rules is morally best (note: for the sake of argument, we can suppose these rules are ‘sufficiently general’):

The Scapegoating Rule: scapegoating is right when the odds of discovery are low enough, and fair trials are wrong.

The Fair Trials Rule: scapegoating is wrong even when the odds of discovery are low enough, and fair trials are always right. (Perl 2021: 99)

According to the Patterned View, to make this evaluation we first need to consider all the good minus bad caused by acts the Scapegoating Rule classifies as morally right (i.e., acts of scapegoating). We then need to consider all the good minus bad caused by acts the Fair Trials rule classifies as morally right (i.e., running fair trials/never scapegoating). The morally best rule is the rule that has the greater good minus bad. (Here, Perl is pretty convinced it is going to be the Fair Trials rule).

The Patterned View is ambiguous in a few respects (Perl might reasonably argue ‘flexible’). Given the analysis to come, two areas are worth highlighting.

One question we might have about the Patterned View is how we ought to read the Generality Requirement. According to one interpretation, we can read the requirement as intended to legislate against only those rules that contain indexicals. Alternatively, we can read the requirement as intended to advantage those rules that contain more inclusive descriptions of the relevant act. Call the former interpretation the *narrow reading* of the Generality Requirement. Call the latter, the *wide reading*.

Perl himself seems to waver between these readings. When he first introduces the requirement, Perl appears to endorse the narrow reading. A comparable demand, he writes, is Rawls’s requirement that ‘principles should be general’, which is to say, ‘it must be possible to formulate them without the use of what would be intuitively recognized as proper names, or rigged definite descriptions’ (cf. Rawls 1971: 113). Later, however, when dealing with a potential counter, Perl seems to endorse something closer to the wide reading. He writes,

The Generality Requirement helps the Patterned View avoid both of these dangers by pushing eligible rules to be more like the rules that appear in a Rossian table of duties. I suggest that rules about dinosaurs in particular violate the Generality Requirement; those rules are not general. Given my suggestion, the Generality Requirement forces us to consider only more general rules, like ‘torturing is right’, or ‘intentionally inflicting gratuitous pain is right’. (Perl 2021: 103)

It may be that in this passage Perl is still looking to endorse a narrow reading of the Generality Requirement. But regardless, the point is this: suppose we think there are at least some instances of ‘intentionally inflicting gratuitous pain’ that do not constitute ‘torture’ (the former is ‘more general’). In that case we face a question: does the Patterned View look to advantage the former rule over the latter? On the narrow reading, the answer is ‘no’. On the wide reading, the answer is ‘yes’. For all Perl says, it is not clear which reading the Patterned View endorses.

The Patterned View is also ambiguous on the question of what agent-neutral good or bad we ought to count in our assessment of the moral value of a given rule. There is a quite staggering range of possibilities here. Given the discussion to come, however, one pertinent issue is the *timeframe* the Patterned View takes to be relevant in that calculation. Here the Patterned View might take into account

- (a) all the extant, agent-neutral good or bad caused by all *past* acts R classifies as morally right.

Or

- (b) all the extant, agent-neutral good or bad caused by all *present* acts R classifies as morally right.

Or

- (c) all the extant, agent-neutral good or bad caused by all *future* acts R classifies as morally right.

Or

- (d) some combination of (a), (b) and (c).

For his own part, Perl seems to opt for a maximalist view under (d). Specifically,

- (d*) the Patterned View ought to take into account all the extant agent-neutral good or bad caused by *all past, present AND future* acts R classifies as morally right.²

At this point, some writers might want to question the coherence of ‘future actual’ good and bad. But I am not really interested in the merits of these different options here. Rather, I flag it only as one area in which the Patterned View appears to yield different possible interpretations, these differences being significant with regards to our following discussion.

Let us now consider the Ideal Worlds objection.

II. The Ideal Worlds objection

Before recounting the substance of the Ideal Worlds objection, it is perhaps worth registering that what *constitutes* the Ideal Worlds objection is, itself, contested. The objection has a long history, having been framed by various writers in various ways.³ Moreover, for obvious reasons, many writers have had a vested interest in interpreting the

²See, e.g. Perl’s claim that ‘this challenge is only pointed if society never switches to the better system. If it does switch at some point, the Consilience Requirement will credit rules enjoining the switch with the future actual good, and thus predicts a requirement to switch’ (Perl 2021: 123).

³See, among others, (Harrison 1952); (Hill 1972: 314–15); (Regan 1980); (DeCew 1983); (Phillips 1985); (Moore 1993); (Korsgaard 1996: 149); (Hooker 2002); (Ridge 2006); (Rosen 2009); (Portmore 2009); (Forschler 2017); (Podgorski 2018); (Miller 2021).

objection differently, as they look to ensure either that some candidate theory is immune to the objection or susceptible to it. At the time of writing, however, I think most writers would agree that the industry standard is Parfit's

The Ideal Worlds objection: Some moral theories incorrectly require us to act in certain ways even when, because some other people are not acting in these ways, our acts would make things go very badly, and for no good reason. (cf. Parfit 2011: 314)

Just as the question of what constitutes the nature of the objection is contested, the question of which theories are open to the objection is also a matter of debate. Since Podgorski (2018), however, philosophers have generally taken the Ideal Worlds objection to apply to all those pattern-dependent moral theories that take moral facts to be grounded in facts about patterns of group behaviour too distant to be activated by any of our actions, or too distant for our actions to place a burden on others to act in certain ways with regards to them.⁴ One such theory is the Patterned View's chief rival, traditional rule consequentialism.

To see the force of the objection against such views, suppose we adopt a traditional rule consequentialist position like

Universalist Rule Consequentialism (URC): We ought to act on only those rules that would make things go best, were everyone to accept them and act upon them when they can.

Now consider the following case

Reluctant: Suppose you and I are contestants on a game show. The show gives us the opportunity to win large amounts of money for a charity of our choice, that will use it to save a significant number of lives. To win the money, all we need to do is coordinate our actions in pressing one of two buttons, Green or Red. The possible outcomes are:

		I	
		Press Green	Press Red
You	Press Green	We save everyone	We save no one
	Press Red	We save no one	We save some people

However, suppose also that, just before we press our buttons, you state truly that you are not going to press Green because you don't want to.

Given your position in *Reluctant*, what should I do? Presumably, our normal intuition here is that I should act on

(A) Press Red, Not Green.

⁴Note: one can see that at least one group this excludes are those pattern-dependent moral theories that take moral facts to be grounded in facts about patterns of group behaviour close enough to be activated by one of our actions (or close enough for one of our actions to place a burden on others to act in certain ways with regards to them). One such theory might be Regan (1980). Another might be Woodard (2019).

Where I endorse (A), we at least save some people. Otherwise, we save no one. However, counterintuitively, URC recommends

(B) Press Green, not Red.

URC recommends (B) because, when choosing between (A) and (B), (B) satisfies its criteria of what makes a rule morally best. (B) is a rule which, if everyone were to accept it, and act upon it when they could, would make things go best. Yet, if I press Green and you do not, we save no one. Thus, as the Ideal Worlds objection claims, URC seems to incorrectly require us to act in certain ways even when, because some other people (i.e. you) are not acting in these ways, our acts would make things go very badly, and for no good reason.

Needless to say, there are various ways in which rule consequentialists and others have attempted to adjust their theories in order to avoid this problem. (Hence why, for example, Hooker's preferred version of rule consequentialism, mentioned above, differs in important respects from URC.) However, Podgorski's (2018) gives us reason to think none of these moves will ultimately succeed. This is because the reason the Ideal Worlds objection has *bite* against pattern-dependent theories is because it picks out a structural weakness in the theories themselves. Specifically, because all these theories claim that what we have reason to do *here and now* is, in some way or other, grounded in facts about patterns of behaviour that are happening *over there* – which is to say, patterns of behaviour too distant to be activated by any of our actions, or too distant for our actions to place a burden on others to act in certain ways with regards to them – they will always be liable to require us to act in ways that, *here and now*, would make things go very badly, and for no good reason. As Podgorski puts it, given this structural weakness, such theories are always liable to miss the 'landmines at their feet'.⁵

III. The Patterned View and the Ideal Worlds objection

We have seen, then, that pattern-dependent views like URC look vulnerable to the Ideal Worlds objection. Our next question is: is the Patterned View? Perl argues 'no'. I think it is.

To see this, let us begin by adopting what I described earlier as Perl's 'maximalist' interpretation of the Consilience Requirement:

(d*) the Patterned View ought to take into account all the extant agent-neutral good or bad caused by *all past, present AND future* acts R classifies as morally right.

Now consider again the case of *Reluctant*. However, suppose that this time we also assume that because the game show is so easy, every previous contestant has always

⁵See related arguments in (Arneson 2005), (Rosen 2009), (Portmore 2009), and (Portmore 2017). Note here: it is not part of my argument in this paper that these arguments, and the Ideal Worlds objection in general, are necessarily unimpeachable. I claim only that this constitutes a standard reading of the objection and, if Perl is to make good on his claim that the Patterned View offers an advance on traditional forms of rule consequentialism with regards to the objection, he must be able to show where the Patterned View is able to avoid the objection where those other theories cannot. Adjudication on that point is compatible with the claim that everyone involved in this debate is wrong to view the objection as, say, a fatal one.

pressed Green, not Red. The possible outcomes of our game, however, remain the same. Namely:

		I	
		Press Green	Press Red
You	Press Green	We save everyone	We save no one
	Press Red	We save no one	We save some people

Finally, as before, suppose also that just before we press our buttons, you state truly that you are not going to press Green because you don't want to. Call this amended version of our earlier case *Reluctant**.

In *Reluctant**, just as in *Reluctant* with regards to URC, it looks as though we ought to endorse

(A) Press Red, not Green.

However, counterintuitively, the Patterned View recommends

(B) Press Green, not Red.

The Patterned View recommends (B) over (A) because, in line with (d*), (B) is credited with all the agent-neutral good caused by actions R classifies as morally right. Since the game show has been running for a while, and since the competitors have saved a lot of people by pressing Green, that is a lot of agent-neutral good and no agent-neutral bad. However, as the Ideal Worlds objection claims, in its endorsement of (B), the Patterned View seems to require me, incorrectly, to act in certain ways even when, because other people are not acting in these ways, my act would make things go very badly, and for no good reason. So, as with other pattern-dependent views, one way we can object to the Patterned View is on grounds of the Ideal Worlds objection.

IV. Potential responses

4.1 Preamble

There are various ways in which a defender of the Patterned View might respond to the argument above. I examine what I take to be their strongest options in a moment. However, before doing so it is important to underline one limit to the kind of counter that such a proponent can offer. As emphasized above, the principal claim that this paper is looking to assess is the claim that one advantage the Patterned View enjoys over numerous other pattern-dependent views – most notably, traditional rule consequentialism – is that, unlike those views, the Patterned View avoids the Ideal Worlds objection. This being so, when assessing potential ways a proponent of the Patterned View might respond to the argument above, we need to make sure we do not allow them to make any moves that are also available to advocates of other pattern-dependent views. Otherwise, those advocates might reasonably claim the Patterned View offers no advantage over their view. In a slogan, we can say

No Specious Arguments: One cannot claim one's view enjoys an advantage over some competitor, where that advantage depends upon one's making an argumentative move that is also available to that competitor.

Let us now consider the possible responses a proponent of the Patterned View might offer.

4.2 ‘Going general’

One way that proponents of the Patterned View might respond to the counterexample offered above is to claim that the best version of the Patterned View is unlikely to recommend rules as specific as (B). Rather, given the Generality Requirement, it would likely recommend much more general rules. Indeed, in his paper, Perl notes one rule he could appeal to here.

The Benefit Rule: choosing a greater benefit over a lesser benefit is right. (Perl 2021: 122)

Assuming the Patterned View recommends the Benefit Rule, it will not recommend that I press Green in *Reluctant**. Pressing Green does not choose the greater benefit over the lesser. Rather, I should press Red.

For a proponent of the Patterned View to make this move is, of course, for them to endorse the *wide* reading of the Generality Requirement noted earlier – that is, one which advantages those rules that opt for more inclusive description of the relevant act. I think this move may cause problems for them elsewhere. However, there are a couple of more immediate problems with this solution.

First, ‘going general’ is a solution available to other pattern-dependent theorists. Both traditional rule consequentialism and, arguably, Scanlonian contractualism have principled reasons for preferring more general rules/principles to more specific ones. Moreover, such theories can be understood to recommend broad duties of beneficence in the style of the Benefit Rule – especially when, as Perl stipulates with regards to the Benefits Rule, such duties are ‘subordinated’ to other rules (Perl 2021: 122). ‘Going general’, therefore, violates *No Specious Arguments*.

Second, ‘going general’ does not solve the problem. Rather, in response we might simply change the case. Consider

Judges: Suppose you and I are judges jointly tasked with sentencing a large group of criminals. If we both adjudicate properly – i.e. in light of the full facts of the case – all the criminals will receive a just sentence. If neither of us adjudicate properly – say, by deciding randomly – the criminals will receive a moderately unjust sentence. But if we disagree on the criminals’ sentence, the decision will go to an Appeals Court, which is famous for handing out extremely unjust sentences. This gives us the possible outcomes:

		I	
		Adjudicate properly	Decide randomly
You	Adjudicate properly	All receive a just sentence	All receive an extremely unjust sentence
	Decide randomly	All receive an extremely unjust sentence	All receive a moderately unjust sentence

However, suppose also that, just before we begin our deliberations, you explain that you’re not going to adjudicate properly, because you do not want to.

Here, again, the Patterned View is open to the Ideal Worlds objection. Presumably here a proponent of the Patterned View will want it to recommend ‘Adjudicate Properly’ (Perl’s chief example of a rule defended by the Patterned View is a very similar ‘Fair Trials’ rule). Yet where I adjudicate properly and you do not, my act would make things go very badly, and for no good reason.⁶

4.3 ‘Biting the bullet’

At this stage, perhaps, a proponent of the Patterned View might simply try to bite the bullet. Here they might remind us of the fact that the Patterned View is a *principled* view. It is not Act Consequentialism. As such, there will necessarily be cases in which it recommends acts that do not make things go best. Nevertheless, such a proponent might argue, it correctly identifies our moral duty. For example, in *Judges I ought* to adjudicate properly, even where you do not.

However, ‘biting the bullet’ is also a move that many other pattern-dependent theorists will want to make in precisely the kinds of cases under discussion. So, it violates *No Specious Arguments*.

4.4 ‘Going conditional’

An alternative way proponents of the Patterned View might look to avoid the counter-examples above might be by employing conditional rules. (Some things Perl says suggest that this is the kind of strategy he thinks proponents of the Patterned View ought to employ.)

We can see how this might work. Take *Reluctant**. In *Reluctant**, a proponent of the Patterned View might claim that we do not have to see the relevant choice as a narrow one between

(A) Press Red, not Green;

and

(B) Press Green, not Red.

Rather, we might see it as a choice between (A), (B) and

(C) Press Green, unless some people are pressing Red, in which case, press Red.

From here they might argue that the Patterned View will not only prefer (B) to (A) but also (C) to (B). First, the Patterned View will credit (C) with the same good it credits (B): the lives saved through successful coordination in pressing Green. However, it will also credit (B) with some bad it does not credit to (C): the lives lost through my dogmatically continuing to press Green when others are pressing Red. So, since the Patterned View will credit the same good to both (B) and (C) but some bad to (B) that it does not to (C), it prefers (C) to (B). And where the Patterned View prefers (C) to (B), its proponents can claim that it is not open to the Ideal Worlds objection:

⁶Note also that here Perl cannot appeal to the Benefits Rule: as he clearly states, the Patterned View will subordinate the Benefit Rule to rules like ‘Adjudicate Properly’ (Perl 2021: 122).

it does not require that I act in certain ways even when, because you are not acting in the same way, my act would make things go very badly, and for no good reason.

Despite this apparent success, however, it seems clear that ‘going conditional’ will not help the Patterned View avoid the Ideal Worlds objection. As traditional rule consequentialists have found, pattern-dependent views of the relevant sort face a series of problems when looking to resolve the Ideal Worlds objection by employing conditional rules. At least some of these problems also affect the Patterned View. Given the nature of the problems in this latter set, we have reason to believe that the Patterned View does not enjoy a significant advantage over traditional forms of rule consequentialism, with regard to avoiding the Ideal Worlds objection.

For example, one problem writers have noted when traditional rule consequentialists ‘go conditional’ is that they can struggle to ensure that their theories will actually adopt the kind of rules that would help them in the relevant range of cases. By tweaking our counterexamples, we can ensure that the consequences of distant patterns of action (patterns that, here and now, we do not have the power to bring about through our own action, nor have the power to place a burden on others to bring about) will push such theories to endorse rules that their proponents would not want them to endorse (say, rules like (B) in *Reluctant**) rather than rules that they would (like (C)).

To see this, take URC and a case like *Reluctant**. Now, suppose that we stipulate that there will be fantastic consequences if everyone obeys (B). (To help us envisage such a scenario, Podgorski (2018) asks us to imagine something called a ‘utility landmine’ – that is, a device that is completely inert until a particular trigger condition is met, where upon, if it is a ‘goodmine’, it will create a near infinite quantity of good things; and if it is a ‘badmine’, it will create a near infinite quantity of bad things. Here, we are imagining there is a goodmine triggered to go off if everyone obeys (B)). In this case, despite proponents of URC hoping their theory might endorse conditional rules like (C), it will prefer (B): in the world in which URC tests its rules – that is, a world of universal adherence – (B) has better consequences than (C). Yet, that brings us back to the Ideal Worlds objection. Once again, URC would wrongly require me to act in certain ways in *Reluctant** (i.e. dogmatically follow (B)), even when, because you are not doing the same, my act would make things go very badly, and for no good reason. Call this the *Misdirection Problem*.

The Misdirection Problem is a significant problem for traditional forms of rule consequentialism when they try to solve the Ideal Worlds objection by ‘going conditional’. However, with respect to that problem, the Patterned View seems to have much the same weaknesses. For example, suppose that, years after we have played our game in *Reluctant**, I invent a badmine that triggers whenever someone obeys (C) and then that, later that week, someone does obey (C) and the badmine is triggered. In that case, the Patterned View would prefer (B) to (C): in the world in which the Patterned View tests its rules – that is, the past, present and future of the actual world – (B) has better consequences than (C). Yet, that brings us back to the Ideal Worlds objection. Once again, in the case of *Reluctant**, the Patterned View would wrongly require me to act in certain ways (i.e. dogmatically follow (B)), even when, because you are not acting in the same way, my act would make things go very badly, and for no good reason.

It may be that a proponent of the Patterned View wants to take issue with the last bit of that sentence – that is the bit that says the Patterned View is directing us to undertake certain actions ‘for no good reason’. We will get to that point in Section 4.5. For the moment, however, we need only register that the *structure* of the problem, and

corresponding objection, is the same. In both cases, the candidate moral theory is directing us to acts which, because others are not acting in the same way (following (B)), will generate harm (saving none where we could have saved some).

In one respect, I do not think we should find this result – namely, that the Patterned View is as susceptible to the Misdirection Problem (and, with it, the Ideal Worlds objection) as traditional forms of rule consequentialism – terribly surprising. This is because, as far as I can see, both the Patterned View and traditional forms of rule consequentialism share the same structural weakness. This is a key respect in which my analysis diverges from Perl's. According to Perl's diagnosis, the reason traditional forms of rule consequentialism are vulnerable to the Ideal Worlds objection is because they test and validate their codes by examining the consequences of *hypothetical* patterns of action (hence, presumably, his desire to eliminate such considerations). If this were true, then the Patterned View *could* avoid the Ideal Worlds objection simply by eliminating all talk of hypotheticals. However, as the analysis above shows, the problem is clearly wider. As indicated in Section 2, what gives the Ideal Worlds objection *bite* against traditional forms of rule consequentialism is *not* that they test and validate their codes by examining the consequences of *hypothetical* patterns of action. Rather, as uncovered by writers like Podgorski and others, it is that they test and validate their codes by examining the consequences of *distant* patterns of action. That is, patterns of action that, here and now, we do not have the power to bring about through our own action, nor place a burden on others to bring about. (This, I take it, is one of the reasons Podgorski strongly encourages us to re-title the Ideal Worlds objection as the *Distant* Worlds objection – or, as I might prefer, the *Distant Patterns* objection). It is this feature of such theories that makes, in Podgorski's phrase, traditional forms of rule consequentialism liable to 'miss the landmines at their feet'. Yet, it is also in this respect that there is no difference between traditional forms of rule consequentialism and the Patterned View. Thus, both are vulnerable to the same objections.⁷

⁷Note: there is no reason to think that this problem is simply generated by our forcing the Patterned View to endorse categorical, as opposed to conditional, rules. That is, we can create the same problem by pushing it to endorse some conditional rules over others. For example, take certain strategies in chess (note: what follows is a toy example, it is not intended to reflect the actual merits of different strategies in chess). Let's suppose that the most successful opening across chess for White is the Italian Game and that winning games of chess generates a small net positive of good over bad. However, let us also suppose that sometimes Black is uncooperative and moves into the Sicilian defence. This gives White two options: employ the Moscow Variation or employ the Rossolimo Variation. The relevant rules here would be:

(D) Play the Italian Game, unless Black plays the Sicilian defence, in which case play the Moscow Variation.

(E) Play the Italian Game, unless Black plays the Sicilian defence, in which case play the Rossolimo Variation.

Now also suppose that, over the history of chess, the Moscow Variation is far more successful than the Rossolimo Variation. In that case, we can imagine the Patterned View will recommend (D) over (E). However, now also suppose that I set up a bot on Twitter that monitors Chess.com and sends huge amounts of annoying spam tweets to every other user every time someone plays the Moscow Variation (I can't stand that variation!). In that case, the Patterned View will now recommend (E) over (D): (D) has consequences we can stipulate to be worse than those realised through adherence to (E). However, note that this also means that, according to the Patterned View, every chess player who has ever played should always play the Rossolimo Variation to the Sicilian defence, despite the fact it will regularly cause them to lose games they could have won by employing the Moscow Variation and despite the fact that the bad consequences realized by *other players* endorsing the Moscow variation would not be realized by *them* doing so: for example, they might be playing centuries before Twitter or Chess.com have even been

Before moving on, let me briefly deal with another point. Earlier I mentioned that there *are* certain problems traditional forms of rule consequentialism face when ‘going conditional’ that the Patterned View avoids. Here is one example. Work by Smith (2010) shows that one problem traditional forms of rule consequentialism face when looking to test codes containing conditional rules is that they struggle to assign them a determinate value. Smith’s examples are perhaps too intricate to rehearse here. However, in brief, the problem arises from the fact that there are multiple ways in which a group of agents may adhere to a code made up of conditional rules and each of these patterns of action can have very different consequences. This might not be an issue if traditional forms of rule consequentialism told us *which* patterns of action obtain in the worlds in which they test their rules. But they do not. They only tell us that, when they test rules, they do so in worlds of such-and-such adherence. On traditional rule consequentialism, therefore, the *expected value* of codes containing conditional rules – the extent to which they tend to result in good or bad consequences – remains indeterminate. We cannot tell whether one code made up of multiple conditional rules is better than another. Call this the *Indeterminacy Problem*.

The Patterned View avoids this problem. Unlike traditional forms of rule consequentialism, the Patterned View *does* tell us which patterns of action obtain in the world in which it tests its rules: namely, that pattern of action that *does* obtain (or *did*, or *will*, or whatever). Thus, on the Patterned View, the *value* of some code will always be determinate (provided that what obtains in the actual world is determinate), even if the code being tested contains multiple conditional rules. There are, then, at least some problems traditional forms of rule consequentialism face when ‘going conditional’ that the Patterned View avoids.

Does this, then, vindicate Perl’s claim that the Patterned View offers a near-Pareto improvement over traditional forms of rule consequentialism? Yes and no. The crucial question here is what we take that claim to imply and, indeed, how ambitious proponents of the Patterned View are about the prospects of their theory.

First, there is a limited sense in which, on the grounds set out above, proponents of the Patterned View can claim the View offers a near-Pareto improvement over other forms of rule consequentialism. Namely, they might claim that it is ‘better on some questions – e.g. the Indeterminacy Problem – but worse on none’.

Now, as I shall show in Section 4.5, I think there are reasons why we might doubt even this claim. (I would argue that the Patterned View *is* better on some questions but also that it is *worse* on some questions.) However, leaving this point aside, as indicated in the Introduction, I take Perl’s claim to be more ambitious than this limited reading of what might constitute a ‘near-Pareto improvement’. That is, I take his claim to be not simply that the Patterned View is ‘better on some questions and worse on none’ but that the Patterned View has an advantage over traditional forms of rule consequentialism in that, unlike those positions, *it avoids the Ideal Worlds objection*. That is a significant claim. Moreover, it is significant not only because it is one problem faced by certain

invented, or playing face-to-face, rather than on Chess.com. This is the Misdirection Problem, paired with the Ideal Worlds objection. Note here also: given the way in which these counterexamples function, we also have good reason to believe that the Patterned View will not be able to escape the Ideal Worlds problem merely by adjusting the Consilience Requirement – say, by endorsing some interpretation of the Consilience requirement other than (d*). (Incidentally, it looks like there will also be significant other problems with that kind of move.)

kinds of rule consequentialists (as is the case with the Indeterminacy Problem). Rather, it is significant because many philosophers consider the Ideal Worlds objection to be a fatal objection to the rule consequentialist's entire project. Were *this* claim to be vindicated, therefore, it would be of interest not only to rule consequentialists but also to *non-rule consequentialists*. It would give *them* reason to sit up and take notice. Yet, as the analysis above shows, even when it 'goes conditional', the Patterned View does not avoid the Ideal Worlds objection. It still sometimes requires us to act in ways that go badly, and for no good reason. Thus, I take it that, with respect to that latter, more ambitious and interesting claim, the Patterned View does not offer us the Pareto-improvement we were looking for.

4.5 Rewriting the Ideal Worlds objection

Let us consider one last way a proponent of the Patterned View might look to respond to the arguments above. One thing they might argue is that it was never their intention to claim that the Patterned View avoids the Ideal Worlds objection in the terms in which we have been discussing it. Indeed, they might even concede that, so described, the Patterned View fares no better with regards to that objection than traditional rule consequentialism. Yet, they might claim that there is another, narrower version of the objection with respect to which the Patterned View does enjoy such an advantage: namely, one specifically addressed to those pattern-dependent theories that ground our reasons for action in facts about hypothetical worlds. Say,

The Ideal Worlds objection (hypothetical worlds): Some moral theories incorrectly require us to act in certain ways *based on a calculation as to what goes on in some hypothetical world* even when, because some other people are not acting in these ways *in the actual world*, our acts would make things go very badly, and for no good reason.

An example can help us make sense of this move. In his (2018), one of the ways Podgorski helps us to see the force of the Ideal Worlds objection is through the idea of a *dud*. That is,

[a utility landmine] of both [good and bad] types with trigger conditions under the following constraint – the trigger conditions are in worlds too distant to be activated by any of your actions, or for your actions to place a burden on others to trigger or avoid triggering them. (Podgorski 2018: 7)

Let us imagine there are some set of duds the trigger conditions of which are restricted to hypothetical worlds/patterns of behaviour, and another set the trigger conditions of which are restricted to actual worlds/patterns of behaviour. Call the former *h-duds* and the latter *a-duds*. From here, a proponent of the Patterned View might claim that, even if the Patterned View is susceptible to a-dud counterexamples (along the lines of, say, *Reluctant**), it avoids h-dud counterexamples (along the lines of, say, the original *Reluctant*).

There are, however, several problems with this move. First and foremost, even if it were successful, such a move would leave uncontested the main thesis of this paper: namely that the Patterned View remains open to a standard interpretation of the Ideal Worlds objection. That would be good news for this paper. But it would be

bad news for proponents of the Patterned View. As described in Section 4.4, many philosophers view susceptibility to *any* dud counterexamples – h-dud or a-dud – to be fatal to a theory’s prospects. After all, it entails that the candidate theory will sometimes require us to make things go badly, and for no good reason. By failing to challenge the idea that the Patterned View is so susceptible, therefore, the proponent of the Patterned View would effectively undercut any motivation any non-rule consequentialists might have for taking an interest in the theory. By their way of thinking, the dispute between the Patterned View and traditional forms of rule consequentialism would be best understood as simply a disagreement between similarly doomed positions.

Second, even if we understand the proponent of the Patterned View’s argument to be simply an argument about which out of (similarly doomed) forms of rule consequentialism is best, it is still not entirely clear why we should think the mere fact that the Patterned View is not susceptible to h-dud counterexamples suggests it has an advantage over more traditional forms of rule consequentialism. After all, in response, traditional rule consequentialists might make a similar claim. They might argue that their position offers an advantage over the Patterned View in that it avoids *a-dud* counterexamples.

In response, a proponent of the Patterned View might attempt to argue something like:

The Asymmetry Thesis: All other things being equal, it is worse for a theory to be susceptible to h-dud counterexamples than a-dud counterexamples.

Some things that Perl says in both his (2021) and elsewhere suggest that this is the kind of argument he wants to make. Yet I am not sure his reasoning here is entirely convincing.

For example, one thing Perl says in his (2021) is that the problem with the way traditional rule consequentialism grounds moral facts in facts about hypothetical patterns of action is *not only* that it means such theories will sometimes direct us to undertake harmful actions but that, when they do, they make such recommendations on the basis of facts that are, in some way, ‘irrelevant’, or something ‘it makes no sense’ for us ‘to care about’.⁸

⁸See, for example, where Perl writes: ‘a core advantage of act consequentialism is its grounding morality in something we can intelligibly care about: our promoting the good. The Patterned View also grounds morality in something we can intelligibly care about. It grounds morality in our reactions to the good everyone does (or, more exactly, in our rational public gratitude or resentment for the good done or omitted). Traditional rule consequentialism is strange compared to either: it grounds morality in things that needn’t ever happen’ (Perl 2021: 124). Elsewhere: ‘the explanation of pacifism [given by, e.g., traditional rule consequentialist theories] seems wrong, that it rests on something irrelevant’ (Perl 2021: 114). As an aside, it is worth noting that there is something a little odd about Perl’s reasoning in the first of these passages. According to Perl, the primary reason why we ought to find rule consequentialism ‘strange’ is that it grounds morality in a counterfactual (something we cannot intelligibly care about). It analyses ‘morally right’ as, e.g., ‘would make things go best, if done by everyone’. However, act consequentialism can also be understood as grounding morality in a counterfactual. It analyses ‘morally right’ as, e.g., ‘would make things go best, if done’. As such, if rule consequentialism is ‘strange’ *qua* counterfactual, act consequentialism must be as well. Perl can thus either maintain that any view which grounds morality in a counterfactual is ‘strange’, in which case he must go back on his claim that act consequentialism ‘grounds morality in something we can intelligibly care about’. Or he can hold on to the idea that act consequentialism ‘grounds morality in something we can intelligibly care about’, in which case he must go back on the idea that rule consequentialism is ‘strange’ *qua* counterfactual. He cannot do both.

One way we can understand this argument is as leaning into the latter part of the Ideal Worlds objection. That is, where that objection argues first, that

some moral theories incorrectly require us to act in certain ways even when, because some other people are not acting in these ways, our acts would make things go very badly,

and second that

When they do so, they do for no good reason,

Perl seems to want to maintain that, irrespective of the susceptibility of either the Patterned View or traditional rule consequentialism to the first part of the objection, traditional rule consequentialism is *more liable* to the second part than is the Patterned View. The thought is that, at least when the Patterned View makes its demands, it does so on the basis of ‘relevant’ considerations, or something ‘it makes sense for us to care about’.

Yet, the operative sense of ‘relevance’ here is obscure. Indeed, on reflection, we might think there is a strong sense in which *both* views make their demands on the basis of something ‘irrelevant’. (So the Ideal Worlds objection is vindicated to that extent.) Namely: both views claim we ought to follow certain rules on the grounds that, in certain circumstances, adherence to such-and-such a rule realizes certain consequences, and that is true even though, here and now, there is *no chance* of those consequences being realized.

For example, take URC. In Section 4.3, we saw that, in cases like *Reluctant**, URC can demand that I follow (B) over (C) because, in the places where URC tests its rules, (B) has better consequences (it sets off a goodmine). Yet, this is even though there is *no chance* of such consequences being realized by *my* following (B) *in that case* (after all, in *Reluctant**, not everyone is following (B)).

The Patterned View has the same problem. In *Reluctant**, the Patterned View can demand that I follow (B) over (C) because, in the places where the Patterned View tests its rules, (B) has better consequences (it doesn’t set off a badmine). Yet, this is even though there is *no chance* of such consequences being realized by *my* following (B) *in that case* (after all, in *Reluctant** *I haven’t even invented the badmine yet*). In this way, then, we might think that irrelevance is *always* going to be a problem for both theories.

Say a proponent of the Patterned View were to concede this. Still, they might argue that there is still a difference between the two cases (one that vindicates the Asymmetry thesis), in that the Patterned View draws its instructions from patterns of action that *will one day happen* (I am, one day, going to set off that badmine). By contrast, traditional rule consequentialism draws its instructions from patterns of action that *may never happen* (it may be that the goodmine is *never* set off).

Yet, again, it is not clear why *this fact* should mean that what *grounds* rules on the Patterned View is something ‘it makes sense for us to care about’ where the same is not true of traditional rule consequentialism. Note, here the proponent of the Patterned View cannot simply argue that, if some rule, R, is to demand an act *for good reason*, it must be the case that the public goods that ultimately justify R *will one day be actual*. To assert such a claim is to beg the question against the traditional rule consequentialist.

Elsewhere in his writings, Perl suggests there may be a way in which proponents of the Patterned View can defend this kind of claim without begging the question. That is, he argues that even traditional rule consequentialists should agree that, if some rule, R, is to demand an act for good reason, it must be the case that the public goods that ultimately justify R will one day be actual, because that follows from their commitment to consequentialism. By virtue of being consequentialists, Perl argues, traditional rule consequentialists should agree that ‘the moral facts supervene on the consequences of agents’ actions’ (Perl 2023: 907). Yet if they agree to *this*, he goes on to intimate, they should also agree that if some rule, R, is to demand an act for good reason, it must be the case that the public goods that ultimately justify R will one day be actual.

However, there is a problem here. We might agree with Perl that any consequentialist must agree that the moral facts supervene on the consequences of agents’ actions. However, we might also think that consequentialists can disagree amongst themselves about whether the moral facts need to supervene on the consequences of agents’ *actual* actions. Some consequentialists might hold that the moral facts supervene on the consequences of agents’ *hypothetical* actions. Such a position is recognizably consequentialist (it is talking about consequences of agents’ actions). Yet it does not entail that if some rule, R, is to demand an act for good reason, it must be the case that the public goods that ultimately justify R will one day be actual.⁹

Here I acknowledge that there may be some readers who agree with all of this and yet who might still think the proponent of the Patterned View has a point. That is, they might think there *is* something odd about the way traditional rule consequentialists ground their rules in patterns of action that may never happen. Yet, in response, traditional rule consequentialists can make an analogous argument. They can argue that there is *also* something odd about the way the Patterned View grounds morality on *only* those patterns of action that *actually* happen.

To see this, consider extramarital affairs. Now, I cannot claim to have any special, first-hand knowledge of extramarital affairs (I promise!). However, it seems plausible to suppose that the overwhelming majority of the agent-neutral bad generated by an extramarital affair is generated upon its discovery. Perhaps there is some bad generated while the affair remains secret. Maybe the unfaithful parties have feelings of remorse and shame. Perhaps the person cheated on has a sense there is something amiss, or is prevented from enjoying certain goods that they might otherwise enjoy (care and attention and so on). But it seems likely that the majority of the hurt and pain associated with such affairs is caused only once the affair is uncovered. If that is correct, and if we also suppose that while they remain secret, such affairs generate a net gain of agent-neutral good (feelings of pleasure and excitement and so on), then it seems possible that the Patterned View will recommend a rule that

⁹Note: this clarification has implications for other aspects of Perl’s argument. For example, one thing Perl argues is that, given the consequentialist is committed to the claim that ‘moral facts supervene on the consequences of agents’ actions’, they are also committed to the following:

Weak Supervenience: Nonnormative duplicates cannot differ in normative respects.

Yet, as the discussion above shows, Weak Supervenience does not follow from that claim. It is possible to be a consequentialist without holding Weak Supervenience. Here I note that Perl suggests some other reasons why consequentialists of all stripes ought to agree with Weak Supervenience. For example, because it is ‘especially uncontroversial’ and because it is ‘compatible with the doctrine of double effect’ (2023: 907). However, given that Weak Supervenience effectively rules out traditional rule consequentialism, it is not clear that these arguments alone give traditional rule consequentialists sufficient reason to endorse it.

(F) Having an extramarital affair that remains undiscovered is right.

Rule (F) will be credited with all the good caused by such affairs. Moreover, since it only takes into account good and bad arising from those affairs that stay secret, it will not be credited with significant amounts of agent-neutral bad. But this looks wrong. Having an extra marital affair is not morally right just because, in the actual world, people often get away with it (something, we might note, recognized by traditional rule consequentialism). Call this the *Getting Away With It Problem*.¹⁰

Even on its own, the Getting Away With It Problem looks like a significant challenge for the Patterned View. However, we might think that it also reflects a deeper, counter-intuitive feature of that position: namely, its unswerving focus on just that possible world that happens to be actual. After all, often what happens, happens, despite what, all other things being equal, we might think *should* have occurred. To put it another way: it often feels like there's no justice in the world – or at least, a lot less justice than we might hope. As such, the world can look like an odd place from which to draw standards of justice. From here, then, traditional rule consequentialists might reasonably claim that the possible world the Patterned View chooses as the grounds for its moral demands looks arbitrary when compared to the hypothetical worlds they consider.¹¹

Overall, therefore, it seems difficult to conclude that, from the mere fact that the Patterned View avoids the Ideal Worlds objection (hypothetical worlds), it has an advantage over more traditional forms of rule consequentialism. And that is leaving to one side the fact that, even if it were established, such a claim does nothing to contest the fact that the Patterned View is open to the actual Ideal Worlds objection.

V. Conclusion

Perl's Patterned View represents a welcome new addition to our existing set of moral theories. Indeed, nothing I have said here need imply that it is not a moral theory worth persevering with. However, with respect to the simple question: 'does the Patterned View avoid the Ideal Worlds objection?' I think the simple answer is, unfortunately, 'no'. With regards to the slightly more nuanced question: 'Does it offer an advantage with respect to the Ideal Worlds objection over other theories, also

¹⁰Proponents of the Patterned View might want to contest this example in another way. For example, they might argue that properly speaking we ought to adopt a wide reading of the Generality Requirement. They might then argue that (F) 'Having an extramarital affair (that remains undiscovered) is right' is insufficiently general. The Patterned View would instead look to assess (G) 'Having an extramarital affair is right' – and presumably reject it in favour of, say, (H) 'Remaining faithful is right'. However, in order for the Patterned View to prefer (H) to (G), it must be the case that, over time, discovered extramarital affairs will cause more bad (minus good) than undiscovered extramarital affairs. Otherwise, the Patterned View will recommend (G) over (H). Thus, the problem still remains: which is to say, as long as enough people have been getting away with it for long enough, the Patterned View will mistakenly classify wrong acts as right.

¹¹Note: it may be that some find the dialectic of this paper frustrating. They might argue that the central question with regards to the merits of the Patterned View over traditional rule consequentialism is not whether it is better placed to respond to the Ideal World objection but rather whether it gives better (more commonsense) advice in a broader range of cases. That is not my interest here. But for those who are interested in that question, the case of 'Having a secret extramarital affair (that remains undiscovered)' seems to me to be at least one case where traditional rule consequentialism gets our moral duties right and the Patterned View gets them wrong.

apparently imperiled by that objection?’ again, I think the answer is ‘no’. Or, at least, ‘not on the evidence we have thus far’. This finding is important, I think, because it helps us better understand the (relative) merits of the Patterned View. However, more significantly, I think it is important because it helps us better understand the extent of the danger posed by the Ideal Worlds objection. It may be that other writers believe they can save their theories from the objection by ‘going actual’ in the way advocated by Perl. This, I think, is unlikely to be successful.

Acknowledgements. For their many helpful comments and suggestions on this and other material on this subject, I would like to thank Chris Woodard and Abelard Podgorski. I would also like to thank the anonymous referees at *Utilitas* and the editor, for a truly excellent set of reviews and suggested revisions.

Competing interests. None.

References

- Arneson, Richard.** 2005. “Sophisticated Rule Consequentialism: Some Simple Objections.” *Philosophical Issues* 15 (1): 235–51.
- DeCew, Judith Wagner.** 1983. “Brandt’s New Defense of Rule Utilitarianism: ‘Ideal Rules and the Motivation to Be Moral’.” *Philosophical Studies* 43 (1): 101–16.
- Forschler, Scott.** 2017. “Universal Practice and Universal Applicability Tests.” *Philosophical Studies* 174: 3041–58.
- Harrison, Jonathan.** 1952. “Utilitarianism, Universalisation, and Our Duty to Be Just.” *Proceedings of the Aristotelian Society* 53: 105–34.
- Hill, Thomas E.** 1972. “The Kingdom of Ends.” In *Proceedings of the Third International Kant Congress*. Dordrecht: Springer, 307–15.
- Hooker, Brad.** 2002. *Ideal Code, Real World: A Rule-Consequentialist Theory of Morality*. Oxford: Oxford University Press.
- Korsgaard, Christine M.** 1996. *The Sources of Normativity*. Cambridge: Cambridge University Press.
- Miller, Dale E.** 2021. “Moral Education and Rule Consequentialism.” *The Philosophical Quarterly* 71 (1): 120–40.
- Miller, Richard B.** 2009. “Actual Rule Utilitarianism.” *The Journal of Philosophy*, 106 (1): 5–28.
- Moore, Andrew.** 1993. “The Utilitarian Ethics of RB Brandt.” *Utilitas* 5 (2): 301–10.
- Parfit, Derek.** 2011. *On What Matters*. Vols. 1 and 2. Oxford: Oxford University Press.
- Perl, Caleb.** 2021. “Solving the Ideal Worlds Problem.” *Ethics* 132 (1): 89–126.
- Perl, Caleb.** 2023. “Some Question-Begging Objections to Rule Consequentialism.” *Australasian Journal of Philosophy* 101 (4): 904–19.
- Phillips, Michael.** 1985. “Reflections on the Transition from Ideal to Non-Ideal Theory.” *Nous* 19 (4): 551–70.
- Podgorski, Abelard.** 2018. “Wouldn’t it be Nice? Moral Rules and Distant Worlds.” *Nous* 52 (2): 279–94.
- Portmore, Douglas.** 2009. “Rule Consequentialism and Irrelevant Others.” *Utilitas* 21 (3): 358–76.
- Portmore, Douglas.** 2017. “Parfit on Reasons and Rule Consequentialism.” In *Reading Parfit: On What Matters*, edited by Simon Kirchin. London: Taylor & Francis, 135–52.
- Rawls, John.** 1971. *A Theory of Justice*. Cambridge, MA: Harvard University Press.
- Regan, Donald H.** 1980. *Utilitarianism and Co-Operation*. Oxford: Oxford University Press.
- Ridge, Michael.** 2006. “Introducing Variable Rate Rule-Consequentialism.” *The Philosophical Quarterly* 56 (223): 242–53.
- Rosen, Gideon.** 2009. “Might Kantian Contractualism be the Supreme Principle of Morality?” *Ratio* 22: 78–97.
- Smith, Holly.** 2010. “Measuring the Consequences of Rules.” *Utilitas* 22 (4): 413–33.
- Woodard, Christopher.** 2019. *Taking Utilitarianism Seriously*. Oxford University Press.

Cite this article: Rumbold, Benedict. 2024. Does the Patterned View Avoid the Ideal Worlds Objection? *Utilitas* 36, 130–147. <https://doi.org/10.1017/S0953820824000037>