

## ARTICLE

# Pattern-Based Reasons and Disaster

Alexander Dietz 

Cardiff University, Cardiff, UK

Email: [alex.dietz@gmail.com](mailto:alex.dietz@gmail.com)

### Abstract

Pattern-based reasons are reasons for action deriving not from the features of our own actions, but from the features of the larger patterns of action in which we might be participating. These reasons might relate to the patterns of action that will actually be carried out, or they might relate to merely hypothetical patterns. In past work, I have argued that accepting merely hypothetical pattern-based reasons, together with a plausible account of how to weigh these reasons, can lead to disastrous consequences. However, in this article, I argue that this problem is not limited to hypothetical pattern-based reasons: it turns out that there are analogous issues for reasons deriving from actual patterns. I then suggest that we can avoid this problem by adopting a different account of the weight of pattern-based reasons.

## 1. Introduction

Many of us feel that our reasons for or against acting in particular ways can depend not only on the features of our own actions, but also on the features of the larger patterns of action in which we would or could be participating: in Christopher Woodard's terms, we feel we have not only *act-based* reasons, but also *pattern-based* reasons.<sup>1</sup> Even though our action might not in itself make a morally significant difference, we may still feel that we should participate in good patterns, and should not participate in bad patterns. For example, you might think that if it would be ideal for us to elect a particular political candidate, then that fact gives you some reason to do your part and vote for that candidate. And you might think that if consumers of factory-farmed meat products are together causing unnecessary animal suffering, then that fact gives you some reason not to participate.

Now, for those who find these thoughts compelling, there are a number of questions to ask. One question is whether one's reasons for action can depend on the part that one would be playing in a merely *hypothetical* pattern of action – an idea prominent in rule consequentialist ethical theories – or only on the part that one is playing in *actual* patterns of action. In particular, suppose that we could together bring about a good outcome, but others will not in fact do their parts. Even though others are not going to join you, do you

---

<sup>1</sup>Woodard has discussed pattern-based reasons in a number of articles and books; a good starting place is Woodard 2008a.

nevertheless have at least some reason to do your part anyway? Another important question is how to measure the *weight* of our pattern-based reasons. How can we decide when our pattern-based reasons are decisive in determining what we ought to do?

In early work, Woodard argued that individuals can indeed have reasons to participate in good patterns even when others will not do their parts.<sup>2</sup> However, in past work, I have argued that when we combine this view with a natural hypothesis about the weight of pattern-based reasons, we can generate unacceptable implications, telling agents that they ought to act in ways that lead to disastrous consequences.<sup>3</sup> In his latest book, Woodard reports that this argument persuaded him to reject the view that we have reasons to participate in merely hypothetical patterns.<sup>4</sup>

I have two aims in this article. My first aim is to show that even if we restrict ourselves to reasons deriving from actual patterns, we still face an analogous challenge. This is because of a simple point noticed by Derek Parfit and Frank Jackson in the 1980s, but neglected in the recent literature: that it can be bad to participate in actual patterns that are themselves good, and can be good to participate in actual patterns that are themselves bad. I show how we can use this point to construct an analogous version of my past argument targeting actual pattern-based reasons.

My second aim is to show that the most promising response to this challenge is to revise our theory of the weight of pattern-based reasons. I first argue that other attempts to defuse the challenge are unpromising. I then examine three possible revisions to the theory of weight. The most plausible option, I argue, is to claim that pattern-based reasons are trumped by sufficiently weighty act-based reasons. This would allow us to maintain both hypothetical and actual pattern-based reasons, while avoiding all of the disastrous implications.

I proceed as follows. In section 2, I review my earlier argument against reasons to participate in merely hypothetical patterns, and in section 3, I show how we can generate similar problems for actual pattern-based reasons. In section 4, I argue against two initial responses to this challenge. Then, in section 5, I propose and defend my account of the weight of pattern-based reasons. Section 6 concludes.

## 2. Hypothetical patterns and disaster

In my earlier article, I argued that when we combine the view that we can have reasons to participate in good but merely hypothetical patterns with a plausible hypothesis about the weight of pattern-based reasons, this can result in disastrous implications. In this section, we will review this argument; in the next section, we will see how this argument bears on our reasons concerning actual patterns.<sup>5</sup>

How can we determine the weight of our pattern-based reasons? I suggested that the weight of pattern-based reasons will depend on how important the pattern is.<sup>6</sup> One way to formulate this idea is as follows:

<sup>2</sup>Woodard 2008a, esp. 81–118.

<sup>3</sup>Dietz 2016: 969–73.

<sup>4</sup>Woodard 2019: 105–8.

<sup>5</sup>The original argument is presented in the context of a discussion of the relationship between the normative reasons possessed by individuals and those possessed by groups of people. However, as we saw above, we may be interested in the idea of pattern-based reasons whether or not we are interested in the idea of group reasons for action. So I will present the argument in terms of the more general idea of pattern-based reasons. This presentation is based on the reconstruction in Woodard 2019: 105–8.

<sup>6</sup>A similar view is proposed in Nefsky 2017: 2764.

*The Proportionality Hypothesis:* The weight of one's reasons (not) to participate in some pattern is proportional to the value (or disvalue) of the pattern, other things being equal.<sup>7</sup>

This hypothesis seems plausible. It also offers an explanation of why our intuitions about the weight of reasons can differ across different cases. For example, on the one hand, we might invoke the idea of reasons to participate in good hypothetical patterns to explain the appeal of maintaining a principled commitment to practices like keeping one's word, even in circumstances in which these practices have broken down: we might feel that we have important reasons to live by the rules of the world we would like to see. On the other hand, consider a case in which two people could paint a house, and it would be good for them to do so, but one of them is not willing to do her part, and the willing painter's contribution would not do any good on its own. In this case, plausibly, if the willing painter still has a reason to do his part, it must be a trivial one. And the Proportionality Hypothesis can explain these judgments: a practice of keeping one's word would be a much more valuable pattern than the painting of a house.

The Proportionality Hypothesis can also explain our varying intuitions about cases involving actual patterns. One example is Parfit's *The Drops of Water*, in which a group of altruists could collect water to alleviate the suffering of men stranded in the desert, but any particular altruist's contribution would only enable each of the men to drink an extra drop, whose effect might be imperceptible.<sup>8</sup> In this case, because the altruists' individual contributions would not make a perceptible difference to any of the men, we may find it hard to see how they could have weighty act-based reasons to contribute. Nevertheless, we might find it plausible that they do have weighty pattern-based reasons to do so, in virtue of the good that they would be doing together, and that these reasons are weighty precisely because so much is at stake. In contrast, it seems plausible that the altruists' reasons to contribute would be much less weighty if the aim were merely to help feed a group of people who had just eaten dinner but were still mildly peckish.

However, if we accept that we can have reasons to participate in merely hypothetical patterns, and we accept the Proportionality Hypothesis as an account of the weight of these reasons, then we can generate counterintuitive implications in certain cases. In particular, there are some cases in which, while it would be great if everyone did their parts in a certain pattern, and not so great if no one did their parts, it would be terrible if some people do their parts while others do not.<sup>9</sup> In these cases, we might think, if other people are not doing their parts, then you should not do your part either; you should not go it alone. Even if we think that we can have reason to do our part in the ideal pattern even if this comes at some cost in the actual world, there is some point at which the cost would become unacceptable.

Now, if we claim that individuals can have pattern-based reasons even when others are not going to do their parts, that does not by itself commit us to saying that you

<sup>7</sup>Woodard notes that my argument does not actually require the idea of a fixed proportion; the argument can also be made with the weaker assumption that the strength of pattern-based reasons increases with the value of the pattern concerned, without approaching an upper bound asymptotically (Woodard 2019: 107). However, I will stick with the proportionality assumption for the sake of simplicity. Other factors that might affect the weight of one's pattern-based reasons might include, for example, how large the causal role one would be playing in the relevant pattern, or what sort of causal role one would be playing. Compare Nefsky 2017: 2764.

<sup>8</sup>Parfit 1984: 76. Parfit's case is based on an earlier case described by Jonathan Glover.

<sup>9</sup>See Gibbard 1965.

should do your part even when that would lead to a terrible outcome. After all, we might say that while you have *some* pattern-based reason to do your part, this can be outweighed by your act-based reasons, and in particular, the facts about the good or bad consequences that your action would bring about. But if we claim that individuals can have pattern-based reasons even when others are not going to do their parts, and the stakes are high enough, then this means that your pattern-based reasons could end up being strong enough to outweigh your act-based reasons not to do things that would lead to terrible outcomes.

To illustrate, I offered an example in which the lives of a very large group of people are in danger, and we face the options shown in [Table 1](#).

**Table 1.** A high-stakes coordination problem

		You	
		do A	do B
I	do A	One hundred saved; the rest die	All die
	do B	All die	All saved

Suppose that, even though we could save the lives of everyone in danger if we both did B, I decide to do A.

In this case, I suggested, it is clear that you ought to do A as well. However, if we accept hypothetical pattern-based reasons, we will have to say that you have a pattern-based reason in favor of doing B, since that is your part in the ideal, though merely hypothetical, pattern that would save everyone. To be sure, we can insist that you also have a weighty act-based reason in favor of doing A, since this action would save one hundred lives. But if we accept the Proportionality Hypothesis, we will have to say that the weight of your pattern-based reason to do B is proportional to the value of the pattern in which we save everyone. And no matter what the relevant proportion is, we can simply increase the number of lives we are supposing to be at stake until the pattern of saving everyone is valuable enough to make your pattern-based reason even weightier than your act-based reason. So we will have to conclude that you ought to do B, even though you will be letting one hundred people die, and saving no one.

Thus, the view that we can have reasons to participate in merely hypothetical patterns, when combined with the Proportionality Hypothesis, seems to generate unacceptable implications. I further argued that it is unclear whether there is any principled way for those who want to claim that we can have reasons to participate in merely hypothetical patterns to avoid this conclusion. I considered and argued against some possible responses to this problem; we will return to these below. In the absence of some alternative solution, I concluded that individuals have pattern-based reasons only when others are willing to do their parts.

This argument may strike some readers as familiar. In fact, this argument can be seen as a variant of the classic “ideal world” objection to rule consequentialism.<sup>10</sup> As

<sup>10</sup>This name for the objection is coined in Parfit 2011: 312–20. Podgorski 2018 argues that a better name would be *the distant world objection*, because he argues that the problematic consequences come from focusing on merely hypothetical or “distant” possibilities, rather than from focusing on overly optimistic possibilities.

we noted earlier, rule consequentialist theories can be understood as claiming that what we ought to do depends solely on whether we would be playing our part in certain ideal hypothetical patterns. According to the ideal world objection, however, these theories can sometimes tell us to act in ways that lead to disastrous consequences, because playing one's part in a pattern that would be ideal if others will also play their parts can be disastrous in a world where others are not going to play their parts.

In his early work, Woodard tried to avoid this problem by developing a pluralistic theory which accepts the rule consequentialist view that we have at least *some* reason to do our part in good hypothetical patterns, but claims that what we ought to do can also depend on the effects of our own individual actions.<sup>11</sup> But my argument suggests that even a more moderate alternative to rule consequentialism which merely claims that we have *reasons* to participate in good hypothetical patterns can imply that we ought to act in ways that lead to disastrous consequences. The only way to avoid the problem, we might think, is to reject reasons to participate in good hypothetical patterns altogether.

Abelard Podgorski has argued that many versions of rule consequentialism can be shown to imply that we ought to act in ways that lead to disastrous consequences, and that this is because these theories make what we ought to do depend on merely hypothetical possibilities.<sup>12</sup> However, Podgorski writes, this does not mean that we need to abandon pattern-based reasons altogether. In particular, Podgorski suggests that we can still accept reasons deriving from objectionable features of patterns that are realized in the *actual* world, or would be realized if we participated; he cites the examples of portraying a fictional character in a way that conforms to an objectionable pattern of stereotyping, and driving a getaway car for a robbery. By contrast, Podgorski writes, standard forms of rule consequentialism focus on reasons for action which “derive from features of social patterns that are *distant*, inaccessible to us, and this is what leads to vulnerability.”<sup>13</sup>

As Podgorski suggests, even if we deny reasons to participate in merely hypothetical patterns, there is still much to be said for preserving the idea of actual pattern-based reasons. In addition, we have seen that there are a number of examples, including those which Podgorski cites, in which the idea of actual pattern-based reasons seems intuitive. Woodard has argued that actual pattern-based reasons also have theoretical attractions; for example, they might be able to provide us with a satisfying explanation of moral constraints.<sup>14</sup> And Conrad Johnson and Richard Miller have argued that there are plausible versions of rule consequentialism that focus on actually accepted ethical rules, rather than ideal hypothetical rules.<sup>15</sup>

As we will now see, however, even if we do reject hypothetical pattern-based reasons, we are still not out of the woods. This is because actual pattern-based reasons, when combined with the Proportionality Hypothesis, can also imply that we ought to act in ways that lead to disastrous consequences.

### 3. Actual patterns and disaster

My argument for the potentially disastrous implications of hypothetical pattern-based reasons relied on cases in which some hypothetical pattern would be (extremely) good, but doing your part, as things in fact are, would be (very) bad. As we have seen,

<sup>11</sup>Woodard 2008a.

<sup>12</sup>Podgorski 2018.

<sup>13</sup>Podgorski 2018: 12.

<sup>14</sup>Woodard 2019: 105–11.

<sup>15</sup>See Johnson 1991 and Miller 2009.

Woodard, Podgorski, and I all concluded from such cases that we should abandon the idea of hypothetical pattern-based reasons, but maintain that we can preserve the idea of actual pattern-based reasons.

However, I will now argue that the assumption that the problem is with hypothetical pattern-based reasons fails to heed the lesson of observations made by Parfit and Jackson in the 1980s.<sup>16</sup> As Parfit and Jackson noticed, there can also be cases in which some actual pattern would be good, but doing your part would be bad, or in which some actual pattern would be bad, but doing your part would be good. I will use these cases to show that actual pattern-based reasons can also generate disastrous implications.

In their discussions, Parfit and Jackson, like Woodard, are interested in how the ethical status of an action might be affected by features of larger patterns of action in which it participates. Unlike Woodard, however, Parfit and Jackson focus exclusively on actual patterns.<sup>17</sup> Parfit initially proposes a principle stating that an act that is not harmful in itself can nevertheless be wrong in virtue of being part of a harmful pattern (such as taking part in a firing squad), and that an act that is not beneficial in itself can nevertheless be required in virtue of being part of a beneficial pattern (as in *The Drops of Water*).<sup>18</sup> However, he and Jackson go on to consider a number of cases that put pressure on these claims.

One of these is a case in which an actual pattern would be good, but doing one's part would be bad. Parfit supposes that an agent has to decide whether to join a group of rescuers in saving one hundred miners. If the agent joins, the rescuers will save the miners, but if the agent does not join, someone else would take the agent's place, and the agent could instead single-handedly save the lives of ten other people.<sup>19</sup> Thus, the agent could participate in a good actual pattern, which would save one hundred lives, but participating would be bad, since this would not itself result in any lives being saved, and it would allow ten additional lives to be lost.

Parfit and Jackson also consider examples of cases in which an actual pattern would be bad, but doing one's part would be good. In particular, they consider examples in which an agent could help to kill someone, but in which the victim would have died anyway, and the agent's participation either spares the victim pain or saves someone else's life.<sup>20</sup> Note that an influential earlier example of a case with this structure is Bernard Williams's Jim case, in which Pedro is planning to kill twenty captives, but offers Jim the chance to help him kill one of them, in which case he will spare the rest.<sup>21</sup>

What is the right thing to do in these cases? Parfit and Jackson suggest that in these cases, the right thing to do may in fact be to refrain from participating in the good pattern, or to willingly participate in the bad pattern.<sup>22</sup> And this seems plausible. Even if you are attracted to the idea that it is often appropriate to participate in good actual patterns, it is reasonable to think that if your participation would lead to a significantly worse outcome, without making things better for anyone, you should not participate. And even if you are inclined to think that it is often inappropriate to participate in

<sup>16</sup>Parfit 1984: 67–86; Jackson 1987.

<sup>17</sup>See Parfit 1986: 848–49.

<sup>18</sup>Parfit 1984: 70.

<sup>19</sup>Parfit 1984: 67–68.

<sup>20</sup>See Parfit 1984: 70–71 and Jackson 1987: 98–99.

<sup>21</sup>Williams 1973: 98–99.

<sup>22</sup>Jackson 1987.

bad actual patterns, it is reasonable to think that if your participation would lead to a significantly better outcome, without making things worse for anyone, you should participate. Parfit and Jackson conclude from these cases that we must either significantly qualify or abandon Parfit's original principle.

I will now argue that these kinds of cases pose a problem not only for Parfit's fairly simple principle, but also for the more sophisticated idea of pattern-based reasons that Woodard went on to develop in the 2000s. Now, as we have seen, Woodard's theory is pluralistic, accepting both pattern- and act-based reasons. We saw that Woodard initially argued that this pluralism gave the theory the resources to avoid traditional worries about the disastrous results of following rule consequentialism. Similarly, the idea that we have reasons in favor of participating in good actual patterns and reasons against participating in bad actual patterns is not in itself incompatible with Parfit and Jackson's judgments about the above cases, since we could still maintain that these reasons would be counterbalanced by our act-based reasons.

The problem, however, is that we can use these cases to generate an analogue of my earlier argument focusing on the weight of pattern-based reasons. In particular, if we accept the Proportionality Hypothesis as an account of the weight of our pattern-based reasons, then we can construct scaled-up versions of the cases in which we will be forced to accept the wrong conclusions. That is, we will be forced to accept that you should participate in extremely good patterns even if your participation would in itself be (merely) very bad, and we will be forced to accept that you should not participate in extremely bad patterns even if your participation would in itself be (merely) very good.

For example, consider the following "bad pattern" case.

A villain is planning to inject two groups of innocent victims with a lethal drug. One group contains a very large number of victims; the other contains one hundred victims. You could administer a second drug to both groups. This drug would have two effects:

- (1) For the members of the hundred-victim group, your drug would save their lives.
- (2) However, because the members of the very large group are allergic to your drug, they would simply die from the combined effects of the two drugs. Administering your drug would not affect the timing or the painfulness of their deaths.

We can provisionally sum up this case as follows: if you participate in the killing of the very large group, then you can save one hundred other people, and if you do not participate, then everyone in both groups will be killed. We will later return to the question of whether administering the drug should really be described as "participating" in the killing of the very large group.

In this case, it is plausible that you have a weighty act-based reason to administer the drug. But suppose we think that you also have a pattern-based reason not to administer the drug, since you would then be participating in killing many people, and that the weight of this reason bears some proportion to the disvalue of the pattern. In that case, as before, no matter what proportion we think this is, we can simply increase the number of the people in the group that you would be helping to kill until the pattern in which you would be participating is disvaluable enough to make your pattern-based reason even weightier than your act-based reason. We will then be forced to conclude that you have overwhelming reasons not to administer the drug, even though this would

mean letting one hundred additional people die, and saving no one. And this, we might think, is the wrong result: in fact, you should administer the drug.

For a “good pattern” case, we could consider a scaled-up version of Parfit’s rescue mission example, in which we suppose that you have to decide whether to join a rescue mission that would save some very large number of lives, but in which your participation is unnecessary, and you could instead save one hundred lives on your own. Similar reasoning would force us to the implausible conclusion that you should participate.

#### 4. Responses

We have seen that, if we accept the Proportionality Hypothesis, then we can find cases where disastrous implications can result not only from hypothetical pattern-based reasons, but also from actual pattern-based reasons. For simplicity, I will focus on the case of participating in an actual bad pattern. In this case, the problem is that we seem forced to conclude that you should not participate in the killing of the very large group, even though your participation would not make things worse for anyone and would save an additional one hundred lives.

In this section, I will argue against two initial strategies for responding to this problem. First, we might say that you would not really count as “participating” in the relevant sense for pattern-based reasons to come into effect.<sup>23</sup> Second, we might insist that the conclusion that you should not participate is in fact correct. I will then defend an alternative solution: accepting an alternative account of the weight of pattern-based reasons.

#### *Would you really be participating?*

Why couldn’t we just say that your action would not count as “participating” in the relevant sense for pattern-based reasons to come into effect? For one thing, it seems to me that your contribution does intuitively count as “helping” to bring about the deaths of the million victims; after all, you are a significant part of what actually causes their deaths.<sup>24</sup> In addition, even if you do not share my intuition about this case, there are structurally analogous cases that can do the same work. For example, I take it that Williams’s Jim example is a fairly paradigmatic case of someone having to decide whether to participate in a morally objectionable course of action. If so, then we could make the argument using a high-stakes version of this case: we can suppose that Pedro offers Jim the opportunity to kill a very large group of captives, in which case Pedro will spare one hundred others.

It’s also worth noting that the action under consideration in the drug example satisfies Julia Nefsky’s prominent account of when participating in a good or bad pattern of behavior plays the right sort of causal role to count as “helping.”<sup>25</sup> Nefsky’s proposal is that if your act of *X*-ing could be part of what causes some outcome, then your act counts as helping to bring about the outcome if and only if, at the time at which you *X*,

<sup>23</sup>For an example of this strategy, see Parfit 1984: 70–73.

<sup>24</sup>Some might deny that something can be part of what causes some outcome if it does not make a difference to whether the outcome occurs. I take it to be common sense that this is possible; for example, that each vote for a victorious candidate is part of what causes the candidate’s victory, even if the candidate wins in a landslide. But in order to address this issue properly, we would need to get into general issues in the metaphysics of causation, which would take us too far afield from the focus of this discussion.

<sup>25</sup>Nefsky 2017.



it is possible that the outcome will fail to come about due, at least in part, to a lack of *X-ing*.<sup>26</sup>

In the drug case, your act of drugging would in fact be part of what causes the deaths of the people in the very large group. And at the time at which you act, it is *possible* that the outcome will fail to come about due to a lack of *X-ing*, because it is *possible* that the villain will have a change of heart and decide not to inject the victims with the lethal drug after all, even if we stipulate that he will not in fact have a change of heart.<sup>27</sup>

Now, one reason why we might be inclined to think that you would not really be “participating” in the killing might be that we have a morally loaded notion of “participation.” In other words, we might be inclined to say that you’re “participating” in a bad pattern of behavior only if we think you’re doing something that you shouldn’t be doing, and so if it’s fine to administer your drug, then you don’t count as participating.

However, pattern-based reasons are supposed to help us to explain why we ought to do what we ought to do, so we need independent criteria for deciding when they apply. If we want to say that what we ought to do in a particular case depends on our pattern-based reasons, then we can’t say that whether we have pattern-based reasons depends on whether we are doing what we ought to do.

Another reason why we might be inclined to think that you would not really be “participating” in the killing of the very large group is that they would die anyway, and your contribution would not in fact make a difference to when or how painfully they will die. However, as we saw earlier, one of the main motivations for pattern-based reasons is that these reasons provide an explanation of why one should or should not participate in some pattern even when it seems that one’s own contribution would not make a difference. So I take it that those who are attracted to pattern-based reasons will be more comfortable allowing that you can count as participating in some pattern even if your contribution would not make a difference.

Perhaps there is some other suitable way of understanding the notion of participation that is able to deliver more acceptable ethical verdicts. But until we find it, we should maintain that you do in fact count as participating in the killing.<sup>28</sup>

We have been considering the idea that perhaps we can block pattern-based reasons from coming into effect by paying attention to whether the agent would be “participating” in the relevant sense. It is worth noting that we might also try adopting other sorts of constraints on which kinds of patterns generate reasons, or in Woodard’s terminology, which patterns are “eligible.” However, whatever kind of pattern we think is eligible, we should be able to reintroduce the problem in the same way I suggested regarding the Jim example: by constructing a structurally analogous case involving the relevant kind of pattern. So this strategy does not look promising either.

<sup>26</sup>Nefsky 2017: 2753. Note that although Nefsky’s proposal involves counterfactuals, that does not mean that we’ve gone back to talking about merely hypothetical pattern-based reasons. Nefsky’s proposal relies on a counterfactual condition to determine when your contribution counts as “helping” to bring about some pattern: you count as helping, on this view, only when the outcome is not guaranteed. But we’re still focusing on what reasons you might have in virtue of the pattern that will in fact come about, not what reasons you might have in virtue of what would happen in some counterfactual scenario.

<sup>27</sup>Nefsky explicitly allows that it can be “possible” in the relevant sense that some outcome will fail to come about because other agents might fail to contribute, even if this is very unlikely. See Nefsky 2017: 2760–64.

<sup>28</sup>In Nefsky 2015 and 2017, Nefsky argues against several alternative proposals for restricting the conditions under which you are to count as “helping” to produce an outcome.

### *Is participating really the right thing to do?*

Let's now turn to a second possible strategy for responding to the problem. Suppose we think that you are in fact participating in the killing. In that case, we might insist that there is in fact nothing wrong with what the argument purported to be the implications of accepting reasons not to participate in actual bad patterns. In fact, we might claim, you do have overwhelming reasons not to administer the drug, notwithstanding the fact that you would not be making things worse for anyone and would be saving a large number of people from harm.

Although there is some plausibility to this position, it seems to me to take an unacceptably strict approach to the idea that we should not participate in harm, one that is in some ways comparable to the absolutist position that we should never infringe deontological prohibitions no matter how much good we could thereby do. After all, if we multiply the numbers of lives that are at stake in our example, we can get the result that you should not participate in the killing even if you would thereby be saving any arbitrarily large number of lives. For example, suppose that if your reason to save  $n$  lives is of strength  $n$ , then your reason not to participate in killing  $n$  people is  $n/1000$ . This means that even if participating in a killing would have the effect of saving one billion lives, we can get the result that you should not participate if you would be helping to kill more than one trillion people.

In their defense, those who take this hardline approach could point out that in order to get this result, we also have to multiply the number of people whom you would be helping to kill. Strictly speaking, they do not need to be committed to absolutism about deontological prohibitions, because they could accept that you should kill or help to kill some number of people in order to save some much larger number, while still maintaining that you should not help to kill a very large number of people in order to save some much smaller, but still large, number. But still, I find it hard to accept that it could be right to let one billion people die if you could save them in a way that does not make things worse for anyone.

Opponents might also go on the attack against my preferred view. They might argue that only those who are already skeptical about constraints against killing in general would think that you ought to participate. However, it seems reasonable to think that while there is a strong constraint against acting in a way that would either kill someone or help to kill someone in cases where your action would shorten this person's life, this constraint either does not apply, or is significantly weakened, when your action would not shorten the person's life. Thus, we could consistently claim both that you ought to participate in the case above, and that in more ordinary cases, we have overwhelming reasons not to perform actions that kill people even if those actions would also save other people's lives.

## **5. Revisiting the weight question**

I have argued that we cannot solve the problem either by claiming that you would not really count as participating in the killing, or by claiming that participating would in fact be the wrong thing to do.

One alternative would be to bite the bullet, and accept that we do not have reasons against participating in actual bad patterns. Many will find this counterintuitive. As we saw earlier, there are a number of reasons why we might find pattern-based reasons attractive. And in particular, we saw that the idea that we have reasons to avoid participating in bad patterns is highly attractive, particularly in contexts where it seems doubtful that one's own contribution will in itself make any significant difference.

Now, to be fair, biting the bullet might turn out to be less costly than it initially seems. It might turn out that in some of these contexts, we will be able to identify act-based reasons against contributing by resisting the initial appearance that your action will not itself make a significant difference, for example, by arguing that there may be some small chance that your action will make a difference, and that the stakes are high enough that your action could still have significant expected disvalue. And while there may still be some remaining range of contexts where this strategy is not successful, these might be cases where biting the bullet seems more acceptable.

However, we also have another alternative. As we have seen, the problematic implications arise when we not only accept pattern-based reasons, but also accept the Proportionality Hypothesis about the weight of these reasons. So if we want to avoid these implications, then rather than abandoning pattern-based reasons, we could instead reassess how to weigh them. In particular, we could adopt a theory of weight which either abandons the Proportionality Hypothesis, or qualifies it with additional principles.

In this section, I will explore three options for how we might revise our theory of weight. The best option, I will argue, is to adopt a rule according to which pattern-based reasons can be trumped by sufficiently weighty act-based reasons. I will then discuss how this issue relates to similar ideas in the literature on rule consequentialism.

### Three options

Here are three options for how we might revise our theory of the weight of pattern-based reasons. First, as I noted in my earlier article, one option is to reject the Proportionality Hypothesis, and instead endorse

*Upper Bound:* There is an upper bound on the weight of pattern-based reasons.<sup>29</sup>

This view could be developed in two ways. On one version, the weight of pattern-based reasons might never reach the upper bound (say weight 10) because it “maxes out” at a particular weight: that is, a pattern-based reason can reach, but never exceed, (say) weight 9.9. On another version, the weight of pattern-based reasons can continue to increase indefinitely, but with diminishing returns, so that it asymptotically approaches the upper bound.

Note also that in order to avoid the disaster problem, we would need to claim that there is an upper bound not merely on the weight of any *particular* pattern-based reason, but also on the total combined weight of an agent’s pattern-based reasons bearing on some action. Otherwise, we could recreate the problem in cases where an agent’s action would count as her part in many different patterns at the same time, since her pattern-based reasons might then sum together to outweigh her act-based reason to prevent the disaster.

A second option is to add the following rule to our theory of weight:

*Simple Trumping:* Other things equal, pattern-based reasons will be trumped by competing act-based reasons.

In other words, we might claim that whenever act-based reasons compete with pattern-based reasons, the pattern-based reasons no longer affect what you ought to do; we just say that the act-based reasons “win.”

<sup>29</sup>Dietz 2016: 972.

Finally, my preferred option, a variant of a view discussed by Woodard, is to endorse a different rule:

*Threshold Trumping:* Other things being equal, pattern-based reasons will be trumped by competing act-based reasons *whose weight meets a certain threshold.*<sup>30</sup>

Let me explain the two trumping accounts in more detail. On these accounts, we have a case of incommensurability.<sup>31</sup> That is, on these accounts, weighing act- and pattern-based reasons cannot be thought of as a matter of simply comparing the values they have on a single cardinal scale. Rather, we can think of act- and pattern-based reasons as measured on two different scales, say, with units  $a$  and  $p$ . On Simple Trumping, in any contest between an act-based reason supporting one action, and any pattern-based reason supporting another action, the agent ought to perform the action supported by the act-based reason. Another way of putting this is to say that act-based reasons will have lexical priority over pattern-based reasons. On Threshold Trumping, act-based reasons will gain lexical priority only when they reach, say, weight  $10a$ . However, along the “pattern” scale of measurement, the weight of a pattern-based reason can keep increasing without limit: the pattern-based reason could be  $5p$ , or  $15p$ , or  $1000p$ , even though it can be defeated by an act-based reason of weight  $10a$  or greater.

Importantly, this means that the trumping accounts are compatible with the Proportionality Hypothesis. That is, our theory of weight could include a rule about when act-based reasons should take priority over pattern-based reasons, while also including the principle that the weight of pattern-based reasons – within their own scale of measurement – is proportional to the value or disvalue of the pattern.

Here is an analogy to illustrate our three options. Suppose that we are designing a card game where hearts stand for act-based reasons, clubs stand for pattern-based reasons, and the rank of a card stands for the weight of a reason. And suppose we want to design the game in a way that ensures that a reasonably high heart – say, a six of hearts – will beat any club that it is played against. If we want to play using the simple rule – “high card wins,” then we could require that the game be played without any clubs above five. Or we could allow players to use a deck with a full set of clubs, but adopt one of two alternative rules: the rule that hearts beat clubs (otherwise, high card wins), or the rule that a six of hearts or higher beats any club (otherwise, high card wins).

Now, we might wonder how much the machinery of the trumping accounts is really doing for us. If pattern-based reasons can be beaten by even the most trivial act-based reason, as Simple Trumping suggests, why not just say that pattern-based reasons have no weight at all? And if any pattern-based reason can be beaten by an act-based reason of a certain weight, as Threshold Trumping suggests, why not just say that the weight of pattern-based reasons has an upper bound? After all, wouldn’t these have equivalent implications about what agents ultimately ought to do?

My first response to this worry is that it is not entirely true that the trumping accounts have the same implications for what agents ought to do as the corresponding accounts mentioned above. In particular, consider cases in which an agent has to choose between two actions, but in which the balance of an agent’s act-based reasons does not

<sup>30</sup>Woodard 2008b. Woodard suggests this view as a possible interpretation of rule consequentialism’s “prevent disaster” rule, which I discuss below.

<sup>31</sup>For an introduction to incommensurability, see Chang 1997. See also Griffin 1986: ch. 5, especially his discussions of “trumping” and “discontinuities” on pp. 83–89.

favor either of them, whether because the agent has no act-based reasons at all (say, because the actions are costless and benefitless), or because her act-based reasons are tied, or are on a par. And suppose that one of the actions would be the agent's part in some valuable pattern. In that case, Simple Trumping would imply that the agent ought to perform that action, whereas the view that pattern-based reasons have no weight at all would not.

In addition, suppose again that the agent's act-based reasons do not support any particular action, but now that the agent has to decide between playing her part in two different patterns, both of which are highly valuable, but one of which is more valuable than the other. In that case, Threshold Trumping will imply that the agent ought to perform her part in the more valuable pattern, whereas the "maxing out" version of Upper Bound may imply that the agent's pattern-based reasons have both maxed out and are therefore tied.

My second response to the worry is simply that we have intuitions about the weights of reasons themselves, and not only about what agents ought to do. As a result, even where two proposals do have the same implications for what agents ought to do, one could be preferable because it better matches our intuitions about the weights of reasons.

### *Evaluating the options*

Let's now assess the costs and benefits of these options. Again, I will argue that Threshold Trumping is ultimately the most promising option.

First, one cost of all three accounts is that they may strike us as ad hoc. That is, even if accepting one of these accounts would enable us to avoid counterintuitive implications, the accounts may not seem particularly natural in themselves, and it is not clear whether we can give them any satisfying rationale. As Woodard writes, what we want is "not mere consistency with the relevant convictions, but an explanation of them."<sup>32</sup>

I grant that this is a significant theoretical cost. However, while we might not be willing to pay this cost merely in order to preserve hypothetical pattern-based reasons, I have argued that the relevant problems also extend to actual pattern-based reasons. And as we have seen, these reasons are also highly intuitive. Once we realize how much is at stake, we might think, accepting some adhocness in our theory of weight is a reasonable price to pay.

Next, a cost borne by two of these accounts, Upper Bound and Threshold Trumping, is that they seem to require drawing an arbitrary line.<sup>33</sup> This line will have to be drawn either where we locate the upper bound, or where we locate the threshold at which act-based reasons have "sufficient" weight. Simple Trumping, by contrast, avoids this cost; rather than endorsing a threshold, it simply claims that act-based reasons quite generally trump pattern-based reasons.

However, note also that it is a familiar idea that we may have to tolerate some arbitrariness in order to strike an acceptable balance between being "principled" and being "pragmatic." As Shelly Kagan points out, anyone who wants to accept a moderate form of deontology, claiming that there are moral constraints against actions such as killing but that these constraints can be overridden when the stakes are high enough, will have to be willing to endorse a seemingly arbitrary threshold.<sup>34</sup>

<sup>32</sup>Woodard 2019: 107. See also Dietz 2016: 972.

<sup>33</sup>See Dietz 2016: 972.

<sup>34</sup>Kagan 1998: 78–81. For a recent discussion of threshold deontology, see Rosenthal 2018.

Turning to the bright side, we have seen that Simple Trumping and Threshold Trumping, in contrast to Upper Bound, are both able to preserve the Proportionality Hypothesis, since they are claims not about the weight of pattern-based reasons themselves but rather about how to settle conflicts between pattern- and act-based reasons. (In terms of our analogy, again, we could still play with a full set of clubs; we would just have to follow one of two rules about hearts taking precedence over clubs.) This is a significant benefit. As we have seen, the Proportionality Hypothesis has significant intuitive support. Indeed, Woodard suggests that a weaker version of the Proportionality Hypothesis, the claim that the weight of one's reasons to participate in a pattern will increase in accordance with the value of the pattern, is likely to be endorsed by any acceptable explanation of the weight of reasons.<sup>35</sup>

The two trumping accounts also enjoy another advantage over Upper Bound. While in many cases, these accounts will all produce the same verdicts about what agents ultimately ought to do, the trumping accounts are better able to capture our intuitions about cases in which pattern-based reasons are competing with each other, rather than with act-based reasons. For example, suppose that I have two options, each of which is supported by a different pattern-based reason. And suppose that both of them relate to highly significant patterns, but one of them is much more significant than the other: say, I must choose between doing my part to end illiteracy and doing my part to avert extreme climate change. Since these patterns are both highly valuable, Upper Bound will likely either say that these reasons have both "maxed out" and are therefore tied, or that one is only marginally weightier than the other, since they would have to be beyond the point of diminishing returns. In contrast, the trumping accounts could imply that the weight of one of these pattern-based reasons could be much greater than that of the other, say  $1000p$  vs.  $500p$ . And it is quite plausible that the trumping accounts would be giving the right answer here. For example, it plausibly makes sense to place much more weight on doing one's part to avert extreme climate change than on doing one's part to end illiteracy.

Unfortunately, Simple Trumping suffers from a very high cost: if pattern-based reasons can be trumped by any competing act-based reasons, then this will undermine much of their practical relevance.<sup>36</sup> This is because pattern-based reasons will then affect what an agent ought to do only in the cases, mentioned earlier, where the balance of act-based reasons does not favor any option. Whenever an agent's act-based reasons do favor one option, even if only slightly, Simple Trumping implies that any conflicting pattern-based reasons will be trumped. In contrast, Upper Bound and Threshold Trumping do allow pattern-based reasons to outweigh act-based reasons outside of high-stakes cases.

Suppose, then, that we rule out Simple Trumping. In that case, I suggest that Threshold Trumping is our best option. While Threshold Trumping suffers from the costs of adhocness and arbitrariness, so does Upper Bound; and Threshold Trumping has the advantages of allowing us to preserve the Proportionality Hypothesis and of capturing our intuitions about cases where pattern-based reasons are competing against each other. The key costs and benefits of the three views are summed up in [Table 2](#).

<sup>35</sup>Woodard 2019: 107–8.

<sup>36</sup>Compare Podgorski 2018: 11–12.

**Table 2.** Summary of costs and benefits

	Upper Bound	Simple Trumping	Threshold Trumping
Avoids adhocness	X	X	X
Avoids arbitrary lines	X	✓	X
Compatible with PH	X	✓	✓
PBR vs. PBR cases	X	✓	✓
Practical relevance	✓	X	✓

Finally, note that the view that I have proposed is not intended to address all potential challenges for pattern-based reasons; for example, the worry that it sometimes seems pointless or futile to do one's part in a good pattern, even when doing so would not come at any significant cost.<sup>37</sup> Threshold Trumping is intended specifically to address the particular challenge that we have been focusing on: the idea that the costs of following these reasons can sometimes just be too high.

### *The prevent disaster rule*

Let's now take a step back. We saw earlier that the original argument against hypothetical pattern-based reasons can be seen as a variant of the traditional "ideal world" objection against rule consequentialism, according to which following rule consequentialism's "ideal code" can sometimes lead to disaster in the real world. But it is also worth comparing Threshold Trumping to another attempt by rule consequentialists to prevent disaster: namely, the "prevent disaster" rule.<sup>38</sup>

The idea is that rule consequentialism will not in fact imply that we ought to act in ways that lead to disaster, and that this is simply because the ideal code would contain a rule (or meta-rule) telling us that we should not follow any of the other rules in any circumstances in which doing so would lead to disaster. In other words, the prevent disaster rule tells us that when the stakes are sufficiently high, our obligation to prevent disaster takes precedence over our ordinary pattern-based obligations. Similarly, I have suggested that sufficiently weighty act-based reasons take precedence over competing pattern-based reasons.

While Threshold Trumping is structurally similar to the prevent disaster rule, though, there are also important differences. First, I have acknowledged that my suggestion does have the theoretical costs of apparent arbitrariness and adhocness, whereas rule consequentialists who endorse a prevent disaster rule do not have to pay those costs. Where do we draw the line for what counts as a "disaster"? As Leonard Kahn points out, rule consequentialists can say that the line is simply whatever line would be drawn by the code that would have the best consequences.<sup>39</sup> And similarly, rule consequentialists can argue that the claim that we should follow the prevent disaster rule is not ad hoc, because it is simply an implication of their basic ethical claim that we should follow whichever code would produce the best consequences, together with

<sup>37</sup>This type of objection has also been raised against rule consequentialism. See, for example, Podgorski 2018's discussion of the "dud factory." For a defense of pattern-based reasons against this objection, see Woodard 2017: 17–18.

<sup>38</sup>Brandt 1992: 87–88, 150–51, 156–57; Hooker 2000: 98–99.

<sup>39</sup>Kahn 2013: 234.

the descriptive claim that a code which included the rule would produce better consequences than one which did not.

Second, however, the appeal to the prevent disaster rule faces a serious challenge. As Podgorski argues, it is possible to construct a scenario in which we can guarantee that the prevent disaster rule will not be part of the ideal code.<sup>40</sup> For example, we can simply stipulate that some machine would produce wonderful consequences if everyone accepted a code that states “Do a jumping jack at noon, no matter what” but not if everyone accepted a code that states “Do a jumping jack at noon, unless doing so would lead to disaster.” And we can further stipulate that if in the actual world, some people do not accept the ideal code, then doing a jumping jack at noon would lead to disaster. In that scenario, it seems that a rule consequentialist would have to claim that you ought to do the jumping jack anyway.

The view that sufficiently weighty act-based reasons trump pattern-based reasons, however, does not face this problem. Again, the prevent disaster rule is supposed to be an implication of rule consequentialism’s basic ethical claim that we ought to follow certain ideal patterns, together with a descriptive claim about what these patterns would look like. Podgorski’s scenario challenges the prevent disaster rule by inviting us to imagine a world in which the descriptive facts turn out differently. But Threshold Trumping is simply offered as an independent ethical principle governing the relationship between act- and pattern-based reasons. This is what makes it vulnerable to arbitrariness and adhocness worries, but this also means that it cannot be invalidated by any inconvenient descriptive facts. That seems to me like a good bargain.

## 6. Conclusion

We have seen that although the idea of pattern-based reasons is attractive, accepting this idea, together with a plausible account of how to weigh these reasons, can sometimes lead to disastrous consequences. And in particular, we have seen that disastrous implications can result not only from the idea of reasons to participate in good but merely hypothetical patterns, but also from the idea of reasons deriving from actual patterns, such as reasons against participating in actual bad patterns. But I have suggested that we can avoid this problem by adopting a different account of the weight of pattern-based reasons: the view that sufficiently weighty act-based reasons trump competing pattern-based reasons. This view could allow us to maintain the idea of pattern-based reasons – for hypothetical as well as actual patterns – without having to accept that we ought to follow these reasons even at very high costs.

## References

- Brandt, Richard B.** 1992. *Morality, Utilitarianism, and Rights* (Cambridge: Cambridge University Press).
- Chang, Ruth,** ed. 1997. *Incommensurability, Incomparability, and Practical Reason* (Cambridge, MA: Harvard University Press).
- Dietz, Alexander.** 2016. What We Together Ought to Do. *Ethics* 126.4: 955–82.
- Gibbard, Allan F.** 1965. Rule-Utilitarianism: Merely an Illusory Alternative? *Australasian Journal of Philosophy* 43.2: 211–20.
- Griffin, James.** 1986. *Well-Being: Its Meaning, Measurement and Moral Importance* (Oxford: Oxford University Press).

<sup>40</sup>Podgorski 2018: 8.



- Hooker, Brad.** 2000. *Ideal Code, Real World: A Rule-Consequentialist Theory of Morality* (Oxford: Oxford University Press).
- Jackson, Frank.** 1987. Group Morality. In J. J. C. Smart, Philip Pettit, Richard Sylvan, and Jean Norman (eds), *Metaphysics and Morality: Essays in Honour of J. J. C. Smart* (Oxford: Blackwell), pp. 91–110.
- Johnson, Conrad D.** 1991. *Moral Legislation: A Legal-Political Model for Indirect Consequentialist Reasoning* (Cambridge: Cambridge University Press).
- Kagan, Shelly.** 1998. *Normative Ethics* (Boulder, CO: Westview Press).
- Kahn, Leonard.** 2013. Rule Consequentialism and Disasters. *Philosophical Studies* 162.2: 219–36.
- Miller, Richard B.** 2009. Actual Rule Utilitarianism. *Journal of Philosophy* 106.1: 5–28.
- Nefsky, Julia.** 2015. Fairness, Participation, and the Real Problem of Collective Harm. *Oxford Studies in Normative Ethics* 5: 245–71.
- Nefsky, Julia.** 2017. How You Can Help, without Making a Difference. *Philosophical Studies* 174.11: 2743–67.
- Parfit, Derek.** 1984. *Reasons and Persons* (Oxford: Oxford University Press).
- Parfit, Derek.** 1986. Comments. *Ethics* 96.4: 832–72.
- Parfit, Derek.** 2011. *On What Matters*, vol. 1 (Oxford: Oxford University Press).
- Podgorski, Abelard.** 2018. Wouldn't It Be Nice? Moral Rules and Distant Worlds. *Noûs* 52.2: 279–94.
- Rosenthal, Chelsea.** 2018. Why Desperate Times (But Only Desperate Times) Call for Consequentialism. In Mark C. Timmons (ed.), *Oxford Studies in Normative Ethics*, Vol. 8 (Oxford: Oxford University Press), pp. 211–35.
- Williams, Bernard.** 1973. A Critique of Utilitarianism. In Bernard Williams and J. J. C. Smart, *Utilitarianism: For and Against* (Cambridge: Cambridge University Press), pp. 75–150.
- Woodard, Christopher.** 2008a. *Reasons, Patterns, and Cooperation* (New York: Routledge).
- Woodard, Christopher.** 2008b. A New Argument Against Rule Consequentialism. *Ethical Theory and Moral Practice* 11.3: 247–61.
- Woodard, Christopher.** 2017. Three Conceptions of Group-Based Reasons. *Journal of Social Ontology* 3.1: 102–27.
- Woodard, Christopher.** 2019. *Taking Utilitarianism Seriously* (Oxford: Oxford University Press).