

Solving the ideal worlds problem*

According to rule consequentialists, moral facts are grounded in facts about the consequences of embedded rules. Paradigmatic embeddings include being widely or universally followed, or accepted, or internalized, by some or most or all of the community.¹ Facts about embeddings of rules are facts about patterns of group behavior – maybe about what happens if everyone accepts the same rule, or if a different moral culture takes hold. Rule consequentialists are thus part of a broader family of moral views – what I’ll call ‘pattern-dependent views’ – where moral facts are grounded in facts about patterns of group behavior. The family also includes contractualists and Kantians. Members of this family differ dramatically from views like act consequentialism, where moral facts are fully grounded in facts about individual actions considered on their own.

Pattern-dependent views face a range of problems, both classic and new, noted in Brandt (1967), Copp (1995), Parfit (2011), Rosen (2009), Podgorski (2018), Lyons (1965), and Smith (2010). I’ll suggest that those problems all arise because of a mistake in standard implementations of pattern-dependence. I’ll then show how to implement pattern-dependence without making the mistake.

Standard implementations rest on *counterfactuals* about how things would go if different patterns took hold. But philosophers have generally grown skeptical that counterfactuals can play this sort of fundamental role, as in Fine (1994, 2001), Martin (1994), and Shope (1978); giving them a fundamental role generates a range of predictable problems. The ambition of this paper is to provide the first counterfactual-free implementation of pattern-dependence. The ultimate argument for what I’m doing will be that it avoids central problems that otherwise arise, including the ideal-world problem and the objection from ‘minimizing conditions’. But even before we see the details, we should *expect* the view to be an improvement over extant alternatives, because it eliminates a device – counterfactuals – that we know causes systematic problems. I’ll suggest that the view is a (near)-Pareto improvement over extant pattern-dependent views. It

*For very helpful comments and discussion, I’m grateful to David Boonin, Heather Demarest, Alex Dietz, Chris Heathwood, Joe Horton, Tanya Kostochka, David O’Brien, Abelard Podgorski, Mark Schroeder, Emily Slome, Julia Staffel, Philip Swenson, Brian Talbot, Jon Wright, the referees and editors at *Ethics*, my 2019 seminar on consequentialism at Boulder, and audiences at a Center for Values and Social Policy talk and the 2020 Eastern. I’m especially grateful to my commentators at the Eastern, Gregory Pence and Teresa Bruno-Niño, to David Copp, who guided me to the best way of thinking about this material, and to Matthias Steup and Garrett Bredeson for material support that let me focus on this project.

¹Richard Brandt (1963), John Harsanyi (1977), Brad Hooker (2000), Derek Parfit (2011), Holly Smith (2010), and Michael Ridge (2006) describe a representative range of the options.

eliminates core problems without losing the advantages that make pattern-dependent views appealing.

1 Pattern-dependence: advantages

Pattern-dependent views are attractive as a way of putting our considered judgments in reflective equilibrium. For example, rule consequentialism promises to do better than act consequentialism in vindicating our considered judgments about killing, promise-keeping, truth-telling, justice, and so on.² Scanlon's contractualism also promises to put our considered judgments in reflective equilibrium in part because it's pattern-dependent.

If we know that we must stand ready to perform actions of a certain kind should they be required, or that we cannot count on being able to perform acts of another kind should we want to, because they are forbidden, these things have important effects on our planning and on the organization of our lives whether or not any occasions of the relevant sort ever actually present themselves. If, for example, I lived in a desert area and were obligated to provide food for strangers in need who came by my house, then I would have to take account of this possibility in my shopping and consumption, whether or not anyone ever asked me for this kind of help (Scanlon 1998, 203).

He takes patterns of expectations to be morally significant even if never executed.

Pattern-dependent views emphasize that agents who coordinate their behavior with others can achieve better outcomes than they could achieve without coordinating; I'll say that they emphasize the *benefits distinctive to cooperative patterns*. Emphasizing those benefits helps explain how morality is something that we can intelligibly care about. We can intelligibly care about the benefits distinctive to cooperative patterns, because we can intelligibly care about what happens to us. So a core advantage of pattern-dependent views like rule consequentialism over rivals like Rossian pluralism is that they give a unified account of what ties all our moral judgments together – benefits of cooperative patterns – and thereby grounding moral facts ultimately in something we can intelligibly care about.

Legal systems are plausible illustrations of benefits distinctive to cooperative patterns. It's best if we all defer to shared legal norms rather than trying to follow whichever legal norms we individually judge to be best. Pattern-dependent views are one of roughly three approaches to benefits distinctive to cooperative patterns. Those benefits might be:

- relevant for what I objectively ought to do even if I can't do anything to bring them about (pattern-dependent views like rule consequentialism or views like Woodard (2008)).

²Of course, each kind of rule consequentialism corresponds to an extensionally equivalent kind of act consequentialism (Portmore 2011; Dreier 1993), which complicates the promise to do better. I won't discuss this complication here.

- relevant for what I objectively ought to do only if I can help bring them about (Regan (1980); Woodard (2019)).
- irrelevant for what I objectively ought to do (act consequentialists).

For act consequentialists, the bare possibility that I and others could bring a benefit about is irrelevant. The possibility matters only when others are *actually* doing their part. And in that case, the benefit is a benefit of my individual action, not distinctive to a cooperative pattern.³ In contrast, a cooperative consequentialist like Regan or later Woodard would take the bare possibility of cooperation to be relevant. For him, benefits distinctive to cooperative patterns matter if others are willing to cooperate.

Rule consequentialists are able to put our considered moral judgments in reflective equilibrium by taking those benefits to be relevant even if we individually can't bring them about.⁴ Consider the sort of schematic case that Judith Jarvis Thomson gives:

You are a sheriff in a small southern town. A murder has been committed, and you do not have the least idea who committed it, but a lynch mob will hang five others if you do not fasten the crime to one individual. (Thomson 1996, 50n2)

We judge that it's wrong for me to scapegoat in this case, even if I can get away with it. Act consequentialists struggle to vindicate this judgment. If I can get away with scapegoating, scapegoating could easily be the action with the best consequences: it results in four more people surviving. This sort of case illustrates how act consequentialists can struggle to put our considered judgments in reflective equilibrium.

In contrast, rule consequentialists are in a much better position to vindicate our judgment that scapegoating is wrong. Their vindication begins by emphasizing that it's very good for a society to have a stable justice system that people trust. People are then willing to have the justice system adjudicate disputes, rather avoiding it because of its capriciousness. And a society where everyone trusts the justice system is much better off than one where people don't. Crucially, though, a trusted justice system isn't the result of any one person's actions. A single sheriff by himself can't collapse trust in the justice system, nor can he repair trust once lost. A trusted justice system is instead the result of a *pattern* of actions.⁵ So if the great good of a trusted justice system is relevant for what the sheriff ought to do, benefits distinctive to cooperative patterns must be relevant even when the agent can't herself bring them about.

³The benefits may still matter in the choice of decision-procedures or dispositions or motives, as discussed later.

⁴As Derek Parfit (2011) and Brad Hooker (2000) both emphasize.

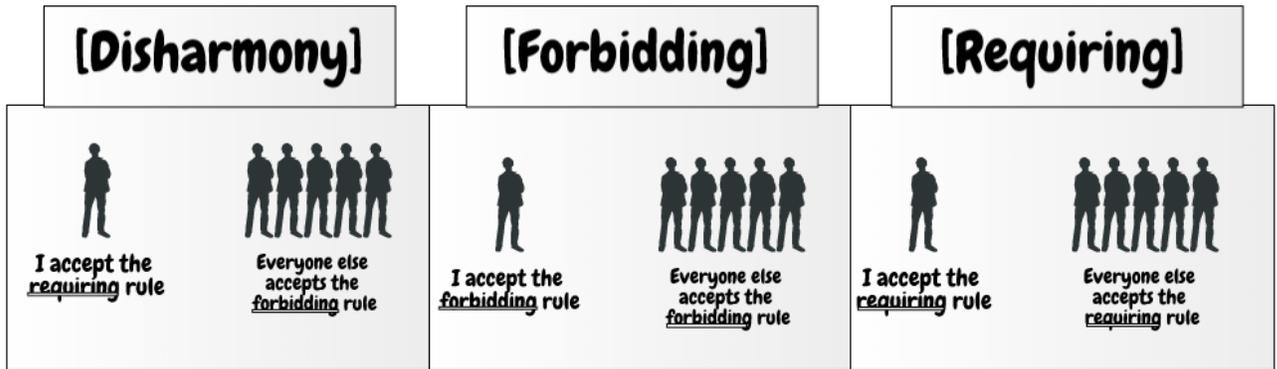
⁵The pattern may cause beliefs about the trustworthiness of the justice system, and those beliefs may be what secures the good. But beliefs don't always mediate the good effects of patterns. So I won't emphasize the mediating role of beliefs in what follows, since I'm interested in the general significance of good patterns.

Rule consequentialist try to capture the moral significance of benefits distinctive to cooperative patterns by contrasting rules about scapegoating:

- The **REQUIRING RULE**: I'm *required* to scapegoat when the odds of discovery are low enough.

The **FORBIDDING RULE**: I'm *forbidden* to scapegoat no matter the odds of discovery.

Rule consequentialists standardly implement their view by contrasting states of affairs:



Let's suppose that:

[Disharmony] is better than [Forbidding], which is better than [Requiring]

This supposition should look plausible. It's a schematic representation of the facts emphasized earlier: that a trusted justice system is a great good, but not one that any one person can cause. Supposing that [Forbidding] is better than [Requiring] represents the fact that a trusted justice system *is* a great good – that good makes [Forbidding] better than [Requiring]. (I assume that enough people would be scapegoating in [Requiring] for trust in the justice system to collapse.) Supposing that [Disharmony] would be even better than [Forbidding] represents two important facts: first, that no one person can destroy by himself destroy trust in the justice system, and second, that my scapegoating would create the good of five people living rather than just one person living.

The good of a trusted justice system illustrates a benefit distinctive to cooperative patterns: a good that's available only if agents coordinate. Since act consequentialists take this kind of benefit to be morally irrelevant, they'll predict that I'm sometimes required to scapegoat. There are only two states of affairs that I can bring about: [Disharmony] and [Forbidding].⁶

⁶For an act consequentialist, we would focus on slightly different states of affairs: ones where I scapegoat or don't. I'm suppressing this complication in the main text, to facilitate the comparison.

Of those two, [Disharmony] is best – and I scapegoat in [Disharmony]. Since act consequentialism tells me to bring about the best state of affairs I can, it'll predict that scapegoating is sometimes required. But predicting that scapegoating is sometimes required is out of sync with our considered moral judgments.⁷

Similar points apply even to the most sophisticated kinds of act consequentialism. Some act consequentialists deny that scapegoating is morally required even if it has the highest expected utility, because they do not focus on the expected utility of acts considered on their own. They instead focus on decision procedures that maximize agents' chances of bringing about what's best. They may insist that the decision procedure that maximizes my chances of doing so forbids acts of scapegoating. After all, scapegoating often has bad consequences, and we're not in a position to know when it does. So you can maximize your chances of doing what's best by refraining from scapegoating. Such act consequentialists are then in a better position to vindicate our considered moral convictions that scapegoating is wrong.⁸

However, even sophisticated act consequentialists still face a problem. Many of us think that scapegoating is *never* morally required. And even though sophisticated act consequentialists can predict that scapegoating *usually* isn't required, it's very hard for them to predict that it's *never* required. Contrast a decision procedure that always forbids scapegoating with a decision procedure that forbids scapegoating unless the stakes are high enough and the odds of discovery low enough. The latter decision procedure could be best. Scapegoating *sometimes* has the best consequences; we should expect the decision procedure that maximizes the chances of doing what's best to require scapegoating when it's likely enough that you can get away with it.

In contrast, pattern-dependent views like rule consequentialism take the great good of a stable justice system into account. For example, a universal acceptance rule consequentialist can predict that scapegoating is always forbidden, because [Disharmony] is simply *irrelevant* for her. Her view is sensitive only to states of affairs where everyone accepts the same rules. Since [Disharmony] involves agents accepting different rules, we put it aside. Between [Forbidding] and [Requiring], [Forbidding] is the only one that contains the great good of a stable justice system, so it is best and determines our obligations: scapegoating is always forbidden.

I've been using the term 'pattern-dependence'; to regiment the use, I'll say that a view that appeals to features of cooperative patterns is (group)-pattern-dependent:

(GROUP)-PATTERN-DEPENDENCE: What's forbidden, required, or permitted is explained by features of a group's pattern of behavior, and not explainable by features of any individual's pattern of action considered on its own.

⁷The same problem arises for Regan's cooperative consequentialism, because his final theory has the property that he calls 'PropAU'.

⁸Moore (1903), Smart (1973), Norcross (1990), Mason (2004), and Railton (1984) have all sketched or defended suggestions with this kind of structure.

Act consequentialism isn't (group)-pattern-dependent, because it focuses on individual acts considered on their own. I'll use 'pattern-dependent' to abbreviate '*group*-pattern-dependent'.⁹

Pattern-dependence has the potential to be a powerful tool in our toolkit. It captures some important moral intuitions in a highly principled way. Because it's such a powerful tool, a wide range of philosophers appeal to it: Kantians and contractualists, as well as rule consequentialists. Unfortunately, though, this powerful tool faces serious problems. The larger ambition of this paper is to get clear on the source of those problems – whether pattern-dependence is itself the source, or if counterfactual elements are the source. So even if you don't fully agree with the previous discussion of scapegoating, take it to be expository – to be an illustration of how someone might think that a pattern-dependent view has important advantages. We'll explore if these views can ever be made to work.

Remember that my ultimate ambition is to provide a near-Pareto improvement of extant pattern-dependent views. The improvement won't solve every problem. We might wonder, for example, if the empirical facts really pattern as smoothly as I've been assuming here. Maybe everyone's scapegoating if they have a 10% chance of detection would be bad. But maybe everyone's scapegoating if they have a 2% chance of detection would be good: the few people who are caught wouldn't be enough to destabilize the good pattern. And in that case, even pattern-dependent views would predict that scapegoating is required in some cases.

But a near-Pareto improvement needn't solve every problem. It's enough if it makes progress to a more defensible view. So I'll bracket further problems by assuming that the empirical facts pattern in the way that would be most helpful for the rule consequentialist – in this case, that universal acceptance even of the 2% rule would be bad. In evaluating a near-Pareto improvement, we need to make sure that the resulting view retains the features of the original view that make it attractive. The empirical assumptions from this section are important for the attractions of rule consequentialism. So we'll assume them in asking whether the resulting view is a near-Pareto improvement.

And even if the crucial empirical assumptions end up false, we can still learn from attempting to reformulate rule consequentialism. My attempt will illustrate a more general strategy for reformulating pattern-dependent views, though I won't develop the more general strategy here. Even if we end up rejecting the crucial empirical assumptions, they can still serve as a ladder to kick away on the route to a defensible pattern-dependent view.

⁹The abbreviation may be misleading; some like Fred Feldman (1975) defend alternatives that center *individual* patterns of action rather than group patterns. But I hope that the context will make the intended meaning of 'pattern-dependent' clear enough.

2 Introducing the Patterned View

The core contribution of this paper is to formulate rule consequentialism freed of *counterfactual* comparisons between possible patterns that may never eventuate. I conjecture that many central problems for pattern-dependent views arise from those counterfactual comparisons. I'll focus on the ideal-world objection that Derek Parfit (2011) described and that Gideon Rosen (2009) and Abelard Podgorski (2018) generalized, together with problems about 'minimizing conditions' that David Lyons (1965) introduced and that Holly Smith (2010) generalized.

We have very general reasons to hope for a counterfactual-free formulation. Though counterfactuals played a fundamental role in mid-century philosophy, many metaphysicians now doubt that they should. Think, for example, of counterfactual analysis of 'fragile' as 'would break if dropped'. We can imagine an angel who'd protect a fragile glass were it ever dropped, with the microstructure of the glass making it fragile even though it *wouldn't* break if dropped (Martin 1994). Robert Shope (1978) gave a general recipe for refuting such counterfactual analysis: roughly, find the cases where the antecedent's happening would affect what it takes for the consequent to happen. Our angel changes what it'd take for the glass to break (the consequent) if the glass is dropped (the antecedent) – the angel then protects the glass. Dissatisfaction with counterfactual analyses feeds into the grounding revolution pioneered by Kit Fine (1994, 2001), Gideon Rosen (2010), and Jonathan Schaffer (2009). (It's no coincidence that Rosen's normative interests make him very concerned about the viability of pattern-dependent views.) And it's the past mid-century that's the historical aberration. As Fine and others emphasize, they're returning to questions that were historically central, not introducing epiphenomena.

Some objections to rule consequentialism implicitly invoke Shope's recipe. One such objection begins by noting that universal acceptance rule consequentialists seem to predict that pacifism is morally required. After all, if everyone universally accepted *don't use violence*, things would go best. But that counterfactual can't figure in a good justification for pacifism! It simply erases the actual violence that might rationalize self-defense. The ideal world objection is that *all* pattern-dependent views issue similarly misguided predictions. This objection illustrates Shope's more general objection: the antecedent (universal acceptance of the pacifist rule) affects the consequent (erasing the actual violence that might rationalize self-defense).

It thus matters if rule consequentialism is hostage to the fate of mid-century analytic metaphysics or if it can be formulated in more contemporary metaphysical settings. As Selim Berker (2018) emphasizes, normative views like rule consequentialism are in part theses in moral metaphysics, so developments in metaphysics can be important for normative work. Now even given the grounding revolution, counterfactuals figure in uncontroversial *heuristics*: fragile things do tend to break if dropped, and treating others as you'd like to be treated tends to matter morally. The problems arise only if counterfactuals have a deeper, non-heuristic role, as in pattern-dependent views like traditional rule consequentialism. Though those sorts

of views in ethics were formulated before the mid-century – say, by Berkeley (1712) – it’s only after the mid-century that we fully came to see the deep and systematic difficulties that arise for foundational use of counterfactuals.

This section shows how to formulate pattern-dependence without counterfactuals by reformulating rule consequentialism. So my account will be an account of how to credit rules as being best, with the best set of rules constituting the demands of morality. Very roughly, the actual good of ϕ -ing will be credited to rules where ϕ -ing is right but not to rules where ϕ -ing is wrong. For example, the actual good of a stable justice system with fair trials will be credited to rules that classify free trials as right and scapegoating as wrong, and not to rules that classify free trials as wrong and scapegoating as right.

In eliminating counterfactuals, this paper reformulates rule consequentialism from the ground up. I do think that others have made important advances in formulating rule consequentialism, and I hope that my reformulation can incorporate many of them. But I’ll focus here just working out the basic reformulation, briefly indicating where I hope it can incorporate those other advances: think of what follows as giving a schematic view that can incorporate additional elements.

2.1 The morally best rules are the demands of morality

My new view begins with the platitude that the demands of morality are the demands of the rules that are morally best. The view will hold, roughly, that the set of morally best rules is the set of sufficiently general rules that has optimal moral value.¹⁰ My development of the platitude incorporates two requirements. The first is the *Generality Requirement*:

GENERALITY REQUIREMENT: the morally best rules
are sufficiently general.

Contrast two formulations of rules about scapegoating:

- *Irreducibly Singular*: If **my** scapegoating is beneficial enough and unlikely to be discovered, then **my** scapegoating is morally right.
- *General*: If **anyone’s** scapegoating is beneficial enough and unlikely to be discovered, then **that person’s** scapegoating is morally right.

The Generality Requirement eliminates the Irreducibly Singular rule. This Requirement starts from our conviction that “because it’s me!” can’t be a brute justification for a moral claim. If it ever looks like a justification, it’s functioning as an elliptical way of referring to general facts about me. A wide range of philosophers accept some version of the Generality Requirement; Rawls, for example, requires that “principles should be general. That is, it must be possible to formulate them without the use of what would be intuitively recognized as proper names, or rigged definite descriptions” (Rawls 1971, 113). I take Rawls to identify paradigmatic *violations* of the

¹⁰I’m grateful to an associate editor for this helpful gloss.

Generality Requirement. But I think he should leave open what else violates this Requirement; for instance, R. M. Hare thinks this sort of requirement should also exclude indexicals (Hare 1952, 129). Whatever Rawls himself thought, I intend to leave the exact contours of the Generality Requirement open for a few pages. Rawls' endorsement at least illustrates the range of philosophers who can endorse a version of this Requirement.

The second requirement gives my new account its consequentialist flair. In formulating it, I'll assume that some states of affairs are good in an agent-neutral way. For example, a state of affairs where I scapegoat an innocent person might be agent-neutrally good, where the benefits of scapegoating outweigh the costs. Act consequentialists can identify moral goodness with agent-neutral goodness: the morally best act produces the greatest amount of agent-neutral good.

The innovation in my new second requirement is to give a different account of the moral significance of what's agent-neutrally good. Roughly, a rule is credited with agent-neutral good/ bad only when the rule "approves" of its production. More carefully:

CONSILIENCE REQUIREMENT: the moral value of a rule \mathcal{R} is everything actual that's agent-neutrally good or bad to the extent it's caused by actions that \mathcal{R} classifies as morally right.

Go back to the earlier example about scapegoating. In the example, my *individual* scapegoating produces something good, but a *pattern* of scapegoating destroys something good – it destroys trust in the justice system.

- 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.

Given the Consilience Requirement, the great good of a trusted justice system is credited to the fair trials rule and is not credited to the scapegoating rule. Fair trials are what sustains trust in the justice system. But according to the scapegoating rule, fair trials are wrong when the person could be scapegoated. So the actions that cause trust in the justice system are wrong according to the scapegoating rule – which means that the great good of a trusted justice system isn't credited to the scapegoating rule. In contrast, the great good of a trusted justice system *is* credited to the fair trials rule. The actions that cause that great good are *right* according to the fair trials rule. Now there is also a genuine good that's credited to the *scapegoating* rule but not to the fair trials rule: the agent-neutral good of successful scapegoating. In the Thomson-style case mentioned earlier, scapegoating would save five people's lives at the cost of only one life. So the agent-neutral good of four more people surviving is credited to the scapegoating rule and not to the fair trials rule. To compare the scapegoating and fair trial rules, we add up the agent-neutral good net bad credited to each.

Our driving observation is that the good of a trusted justice system is greater than all the goods that come from successful scapegoating.¹¹ We're looking for an account of the moral significance of our driving observation free of counterfactuals. That's *exactly* what the Consilience Requirement provides: it predicts that the fair trials rule is best given that driving observation. The Consilience Requirement thus predicts that fair trials are right and scapegoating is wrong, because that's what the best rules say.

I'll call the resulting view the Patterned View.¹² The Patterned View generates a bipartite division: actions are either right or wrong. The more familiar tripartite division of actions as required, permitted, or forbidden is a derivative consequence of the bipartite division. An action is morally required when it's the *only* right action, and an action is merely permitted when it's one of several right actions. (It's easiest to formulate the Patterned View as generating the bipartite division first and the tripartite division derivatively.)

I've contrasted a rule where scapegoating is right with a rule where fair trials are right. The Patterned View does not consider a rule where *not* scapegoating is right, because there are many ways to not scapegoat; for instance, shouting at Reviewer Two is a way of not scapegoating. And the merits of shouting at Reviewer Two are irrelevant for scapegoating. For that reason, the Patterned View bans rules with descriptions of actions that incorporate logical connectives ('not', 'or', ...): there are irrelevant ways of doing actions describable with those connectives.¹³ There are several further technical questions about the Patterned View; for example, whether it requires that the eligible rules classify *all* actions as right or wrong, or whether some sets of eligible rules can be incomplete. Think of this paper as motivating the interest of that sort of further technical question. This paper's goal is to convince you that some version of the Patterned View is well worth taking seriously, so that you are interested in the next paper that works out the technical details of the Patterned View.

The Patterned View has two key advantages. First, it is genuinely pattern-dependent: right and wrong are grounded in facts about the value of patterns, rather than the value of actions considered on their own. We know that it's pattern-dependent because it makes the prediction about the wrong of scapegoating that distinguishes pattern-dependent views from their rivals. Second, it is free of anything counterfactual. Inasmuch as there are problems distinctive to the use of counterfactual elements, the Patterned View avoids them – as §4 will describe in more detail.

Now causation does play a central role in the Patterned View. And some philosophers favor counterfactual analyses of causation. If you favor a

¹¹That observation is necessary for rule consequentialism to be a genuine rival to act consequentialism. And as §1 emphasized, we're here assuming that it is a genuine rival in the course of trying to find the best formulation.

¹²Christopher Woodard (2008, 2019) similarly emphasizes 'pattern-based' reasons. Even though my view works dramatically differently than his, 'pattern/ed' is such a good word for the phenomenon that I'm echoing his use.

¹³If you wish, you may think of this ban as a further Requirement incorporated into the Patterned View.

counterfactual analysis of causation, you can substitute it in here. However, there seems to be a tension between the initial motivation for the Patterned View – formulating pattern-dependent views without counterfactuals – and incorporating a counterfactual analysis of causation. The tension is that someone who accepts a counterfactual analysis of causation will think that the problems arising from the use of counterfactuals can sometimes be surmounted.

But the tension is in the end superficial. Philosophers who accept a counterfactual analysis of causation should think that *something special* about causation allows for a counterfactual analysis; they don't think that their counterfactual analysis hangs on the plausibility of counterfactual analyses of other important relations or properties. So they should admit that it's very much looking for counterfactual-free accounts in other domains, because of the systematic problems that arise for the use of counterfactuals *elsewhere*.¹⁴ Furthermore, I think there is decisive evidence that systematic problems *do* arise for traditional rule consequentialists, whether or not counterfactual accounts of causation succeed. I describe the evidence later, starting in §4.1. So at this point, I haven't introduced the problem I want to solve in detail. I'm trusting you to anticipate the sorts of problems that traditional rule consequentialism face.

Now because the Patterned View is still so schematic, it can incorporate the advances that others have made in formulating pattern-dependent views. We noted Scanlon's emphasis on the importance of being disposed to give aid, and rule consequentialists sometimes emphasize similar dispositions. The Patterned View can incorporate rules about dispositions as well as about actions: credit rules about dispositions with the agent-neutral good due to the disposition.¹⁵ Another important advance in formulating pattern-dependence has been broadening the embeddings that matter – for example, focusing on *internalized* rules, rather than just behavior that follows the rules. My initial formulation of the Patterned View in effect focuses on the consequences of individuals' *following* the rules: it credits the rules with the good that individuals produce. But extensions of the Patterned View could incorporate more sophisticated options. For instance, one extension would credit a rule only with good produced by those who've internalized the rule, or those who act from a recognition of the reasons

¹⁴There's another reason to take the tension to be superficial. Contemporary counterfactual accounts like Judea Pearl (2000) and James Woodward (2003) don't take counterfactuals to figure in fundamental explanatory grounds. But as Selim Berker (2019) emphasizes, normative theorizing aims at identifying the fundamental explanatory grounds. So incorporating a Pearl/Woodward-style view in the Patterned View still wouldn't mean that counterfactuals play a fundamental role.

¹⁵Such rules occasion interesting further questions; for example, what to do if the best dispositions give different instructions than the best rules focused just on actions. The important point here is that the Patterned View allows for those questions to be formulated in much the way they can be formulated in other settings.

that the rule takes to matter.¹⁶

Strikingly, though, the Patterned View is pattern-dependent even before we introduce the sophisticated extensions: even without them, it's genuinely distinct from act consequentialism. In contrast, traditional formulations of rule consequentialism *need* the sophisticated extensions to avoid collapse into act consequentialism. Remember that I'm trying to rebuild rule consequentialism from the ground up, while leaving further extensions open. I find it promising that my reformulation works even given the simplistic option that dooms traditional formulations of rule consequentialism. We can reasonably be hopeful that my reformulation will work even better given more sophisticated options.

2.2 Generality and the poverty of grounds

The Generality Requirement plays a crucial role in the Patterned View that I haven't highlighted yet. Consider the rule 'torturing dinosaurs is wrong'. Given only the *Consilience* Requirement, that rule is in itself no better than the rule 'torturing dinosaurs is right'. Since nobody was around to torture dinosaurs, neither rule could be credited actual good or actual bad. It thus looks like our world lacks sufficient grounds for the wrong of dinosaur torture. In contrast, traditional rule consequentialists might allow for the right sorts of grounds: counterfactual worlds where humans and dinosaurs existed simultaneously.

A related danger is that the Consilience Requirement seems to have a status quo bias: that it can only ratify roughly the moral judgments that we already make. It seems to have a status quo bias, since the only goods that matter for the view are actual goods.

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'. Those rules will be credited with the bad effects of torture, while relevant alternatives like 'torturing is wrong' won't be. So the latter rule will be better. And the latter rule classifies dinosaur torture as wrong, because it's a species of torture, even absent witnesses to the wrongness of torturing dinosaurs in particular. In fact, I see this explanation as a *virtue* of the Patterned View. It forces us to focus on the general features that matter, rather than considering silly possibilities where humans and dinosaurs co-exist as witnesses to the wrongness of dinosaur torture.

¹⁶Such extension incurs commitments about our actual psychology – for instance, that people actually act from internalized rules. Traditional kinds of rule consequentialism don't, because they consider counterfactual scenarios where people are guaranteed to have the psychological states that matter according to the view. So the Patterned View can incorporate extensions about internalized rules only given sufficiently minimal construals of acting from internalized rules. I briefly introduce my favored extension at the end of §3.

The ambition of the Patterned View is thus to put our considered judgments in reflective equilibrium by a *combination* of the Consilience Requirement and the Generality Requirement. The hope is that the Generality Requirement guarantees that sufficiently general rules are the only eligible rules, and the actual grounds can discriminate among the sufficiently general rules. For example, the Patterned View may well support some suitably refined versions of Ross' table of *prima facie* duties. If it does, it's because the Generality Requirement guarantees that eligible rules have the same generality as those in that Rossian table, and the Consilience Requirement ends up predicting that those rules are best.

I introduced the Generality Requirement as articulating a widely shared consensus – that “because it's me” isn't a brute moral justification.

GENERALITY REQUIREMENT: the morally best rules
are sufficiently general.

I take the Patterned View to be highly principled in part because the Generality Requirement articulates a widely shared consensus, which Rawls and Hare both defend. I've added the suggestion that the Requirement excludes rules about dinosaurs in particular. But that suggestion goes beyond the widely shared consensus. For one thing, excluding the concept DINOSAUR as insufficiently general likely means excluding the concept HUMAN as insufficiently general too, since both concepts pick individuals out by singular reference to a particular species. And there's less consensus about whether HUMAN is sufficiently general – that's why arguments about speciesism are intelligible.

I thus distinguish two versions of the Generality Requirement: an initial, unrefined formulation and a final, fully refined formulation. The unrefined formulation articulates the widely shared but incomplete consensus, while the final, fully refined formulation goes beyond it. The *final* formulation of the Generality Requirement needs to settle exactly which concepts are sufficiently general; for example, it needs to settle whether DINOSAUR and HUMAN are sufficiently general. But the final formulation needn't be the object of a widely shared consensus. It's instead justified by the method of equilibrium: it's justified as what combines with the Consilience Requirement to put our considered judgments in reflective equilibrium.

But I've hyped the Patterned View as genuinely-patterned dependent. The hype is right even given only the widely shared but incomplete consensus. That consensus determinately excludes rules with proper names and indexicals – for instance, rules about what I myself may do. That exclusion *by itself* makes the Patterned View pattern-dependent. For instance, the consensus excludes the rule that I alone may scapegoat while others must hold fair trials. Without that exclusion, the Consilient Requirement would credit that *indexical* rule as better than just a rule forbidding scapegoating for absolutely everyone. (It'd credit it with the good consequences of my scapegoating *together with* the great good of a stable justice system.) But *with* that exclusion, the best rule requires fair trials. The advertised advantages of Patterned View rest only a modest and widely shared consensus about which concepts are general. Those advantages are secure even if

the consensus doesn't settle whether HUMAN and DINOSAUR are sufficiently general, and even if further problems arise if they're insufficiently general.

We should only expect a full account of the Generality Requirement at the close of inquiry, not the outset. Moreover, the final formulation of the Generality Requirement need not strike non-consequentialists as facially plausible, since it's ultimately justified by its role in the Patterned View. We should allow for revision in our initial judgments in light of principles like the Generality and Consilience Requirements, while refining those Requirements in light of our initial judgments; Rawls emphasizes that the method sometimes involves "withdrawing our judgments and conforming them to principle" (Rawls 1971, 18). But we don't need the final statement of the Generality Requirement to see the *initial* formulation of the Generality Requirement as a highly principled element of the Patterned View. In other words: it's not an *ad hoc* element of the view, because it incorporates a widely shared if incomplete consensus. And that consensus is enough *by itself* to secure the advantages of pattern-dependence that I've hyped.

The Generality Requirement also eliminates many forms of status quo bias. It forces us to identify the general features of valuable patterns, and to treat those considerations as having the same moral significance wherever they occur. Consider, for example, the line of reasoning from Singer (1972) that concludes that we ought to relieve suffering everywhere: there is no morally significant difference between letting a child drown in front of you and letting someone abroad die when you could save either. To press the status quo objection, suppose that no one actually relieves suffering abroad. Singer's argument is still compelling giving that supposition – for him, it doesn't matter what people are actually doing. But the Patterned View focuses only on *actual* good, so the supposition might make a difference for the Patterned View.

However, the Patterned View still vindicates Singer's line of reasoning, even if no one actually relieves suffering abroad. Consider these three rules:

(*Singer's Rule*) "If it is in our power to prevent [someone from dying], without thereby sacrificing anything of comparable moral importance," (Singer 1972, 132) doing so is right, and letting them die is wrong.

(*Nearby Rule*): letting someone in your immediate environment die is wrong, but letting someone further away die is right.

(*Apathy Rule*): letting anyone die is always right.

To formulate the challenge to the Patterned View, we're supposing that people only save those nearby. The Patterned View credits the good of saving those nearby to Singer's Rule and the Nearby Rule both, but not to the Apathy Rule.

Crucially, though, the Generality Requirement *eliminates* the Nearby Rule as insufficiently general: we recognize 'immediate' as an indexical. Singer himself appeals to a version of that Generality Requirement:

I do not think I need to say much in defense of the refusal to take proximity and distance into account. The fact that a person is physically near to us, so that we have personal contact with him, may make it more likely that we shall assist him, but this does not show that we ought to help him rather than another who happens to be further away (Singer 1972, 232)

If Singer is right, the Patterned View vindicates Singer's Rule as best. Even the initial, unrefined version of the Generality Requirement eliminates the Nearby Rule – unsurprisingly, given Hare's influence on Singer.¹⁷ And the Consilience Requirement makes Singer's Rule better than the Apathy Rule, since it only credits the former with the actual good of saving those nearby.¹⁸

Now you may disagree with Singer – you may think that the Nearby Rule survives his objection. But in that case, *you wouldn't think that the Patterned View delivers the incorrect verdict about Singer's argument.* You'd think that the Nearby Rule is correct; you wouldn't see a failure to vindicate Singer's Rule as objectionable. The crucial point is that someone who finds Singer's argument compelling will also think that his argument succeeds if the Patterned View is true, and furthermore that it succeeds even if no one is *actually* saving those far away. The Generality Requirement thus gives the Patterned View powerful resources for avoiding status quo biases: it allows those who think that Singer's argument succeeds to continue thinking that it succeeds if they accept the Patterned View. There's more to say about status quo biases; I'll continue describing the power and reach of the View's resources as the paper goes on.

¹⁷I expect that the final, fully refined version of the Generality Requirement also excludes rules where our obligations only extend to *children* who we could save from death. Singer presumably sees being a child in the same way he sees being physically near: features that might matter by affecting what it is in our power or in their power, but not features that are in themselves morally significant. Alternatively, though, we might also find that some considered moral judgments *require* rules about children in particular. We'd then learn that the final version of the Generality Requirement allows rules about children in particular. The Patterned View would then interpret Singer's argument as only vindicating obligations to far-off *children*. But the important point remains true: the Patterned View allows for radical moral requirements even if no one is actually acting according to those requirements.

¹⁸The Patterned View thus faces standard 'demandingness' objections – which it should! I introduced it as a near-Pareto improvement on extant kinds of rule consequentialism, which are also demanding. And the Generality Requirement will give the Patterned View a radically different structure from the sort of actual rule consequentialism that Richard Miller (2009) or Conrad Johnson (1991) defend.

3 On rule worship

The Patterned View is itself an original contribution; no one else has defended anything similar. It captures the advantages of pattern-dependent views in a very new way. As later sections will argue, I think it cleanly solves a range of core problems. But even if you're not fully convinced that it cleanly solves those further problems, it's important to appreciate the range of ways that pattern-dependent views can be true. And one of this paper's core contributions is to expand our understanding of that range.

Now a standard complaint about rule consequentialism is that it involves a kind of 'rule worship' – a fetishistic concern for rules, rather an intelligible focus on the good that could be done. This section introduces a new answer to this complaint, then uses that new answer to support the Patterned View. I want to describe this distinctive answer before displaying the advantages of the Patterned View, to convince you that it's a highly principled way to secure those advantages.

To make the rule worship complaint more precise, start with a set of rules identified by some counterfactual profile. Maybe, for example, they're the rules whose universal acceptance would make things go best. Rosen complains "if [someone] tells me that he cares about conforming to such rules for the sake of conforming to such rules, then I would probably not believe him. I don't see how it could make sense to care about such conformity for its own sake" (Rosen 2009, 96). There is something highly intuitive about this complaint. Morality seems like something we can intelligibly care about, but rules distinguished only by some particular counterfactual profile don't.

3.1 Background assumptions

This subsection introduces two background assumptions essential to my answer to the rule worship complaint.

The first background assumption is that there is a distinction between two levels of moral thinking. There is an *ordinary* level of moral thinking. At this level, I might reason from something like Ross' table of prima facie duties– that is, from a fairly settled body of convictions about what we ought to do. But there is also a *critical* level of moral thinking, where I evaluate if my settled body of convictions is correct or mistaken. I'll take rule consequentialists to be offering an account of the critical level, and Scanlon's contractualist to be offering a competing account. When beliefs at the ordinary level cannot be defended at the critical level, we should change them.¹⁹

We should recognize the two levels because we should recognize that moral reasoning plays two different psychological roles. The normative ethics seminar room exemplifies one role: thinking in a cool hour about our moral obligations. But we also reason about morality when our time

¹⁹I intend the metaphor of levels as an homage to R. M. Hare (1981), who developed a moral theory with a similar, two-level structure. But my metaphor is an homage, not a copy; there are significant differences.

or interests are limited. I might be faced with a choice between finishing comments on my student’s paper and calling my mother on her birthday. In thinking about this choice, I weigh the different moral considerations that bear on the decision, but I don’t work the choice out from first principles. When my time or interests are limited, I’m reasoning at the ordinary level, and when I’m in the seminar room, I’m reasoning at the critical level.

This distinction is worth caring about. Eliminating the ordinary level would be objectionably burdensome; ordinary ethical reasoning would take too long. It’d be working things out from first principles, each time. Eliminating the critical level would be objectionably conservative, because defenders of slavery would have no occasion for critical reflection. I’ll assume that the ordinary level consists of *rules* that are selected at the critical level. Take this assumption to be expository: “rules” is just my label for whatever sort of thing we select at the critical level. We could also call them “decision procedures”. It’s intelligible to care about the distinction between the two levels of moral thinking, because of their different roles.

In fact, I take the distinction between the two levels of moral thinking to be largely uncontroversial. It’s the same sort of distinction that act consequentialists like Railton (1984) draw between relying on decision procedures or rules of thumb and the sort of reasoning we do when we’re working things out from first principles. Railton (1984) and Parfit (1984) have both argued compellingly that this distinction is at least *rationally* defensible. Now it *is* controversial whether moral reasoning should look different at the ordinary and critical levels; critics of ‘esoteric’ morality hold that it shouldn’t. The most plausible criticisms allow that the distinction is in principle rationally defensible but insist that features distinctive to morality force moral reasoning at the two levels to look the same.²⁰

My second background assumption will be, roughly, that critics of esoteric morality are right. I’ll call this assumption my *Publicity* assumption:

Publicity: the two levels of moral thinking are both perspectives on our objective moral reasons.

Publicity is controversial. Act consequentialists like Railton think of the ordinary level of moral thinking as consisting of rules of thumb that fail to capture our objective moral reasons. Though it is intelligible to reject this supposition, I think it’s also intelligible to accept it. Accepting it is a way of seeing morality as constitutively *public*. When we give our moral reasons to each other in ordinary life, we’re giving each other the objective reasons that justify our actions, not giving rules of thumb that we aim to follow.

To illustrate, suppose that a doctor can cut up one healthy innocent person to save five sick people. Most of us think that cutting up the one is wrong. Act consequentialists like Railton can *agree* that it’s wrong – at least, they can agree that we should all believe that it’s *subjectively* wrong, wrong given the evidence that we have. For them, the fact he’s an innocent person is a decisive *subjective* reason to not cut him up. Because his

²⁰See Katarzyna De Lazari-Radek and Peter Singer (2010), Hooker (2016), and Singer and De Lazari-Radek (2016) for recent illuminating discussion.

innocence is a decisive subjective reason, we should always cite it when explaining to others why we didn't cut him up. Now as act consequentialists, they hold that cutting him up is sometimes *objectively* right. So they deny that his innocence is always a decisive *objective* reason to not cut him up. As a result, they claim that we know that we should sometimes explain our actions to others by appeal to reasons that are not our objective moral reasons. That claim is a violation of Publicity.

I claim that we can intelligibly care about Publicity without objectionable fetishism. I'll show that Publicity leads to the Patterned View. Since the motivation for Publicity is nonfetishistic, accepting the resulting view won't be fetishistic either. As a result, I don't expect you to find my second background assumption – Publicity – uncontroversial in the way I do expect you to see the first assumption as uncontroversial. Publicity plays a different role in what follows: providing nonfetishistic motivation.

3.2 How Publicity leads to the Patterned View

As Scanlon emphasizes in 'Contractualism and Utilitarianism', normative theories are especially attractive when they ground moral facts in considerations that we intelligibly care about. Focusing on classic utilitarianism, he notes "it seems evident to people that there is such a thing as individuals being made better or worse off. Such facts have obvious motivational force; it is quite understandable that people would be moved by them in much the same way that they are supposed to be moved by moral considerations" (Scanlon 1982, 107). For Scanlon's utilitarian, caring about morality is intelligible because caring about welfare is intelligible. The Patterned View agrees. Caring about morality is intelligible because caring about agent-neutral good more generally is intelligible.²¹

But caring about agent-neutrally good isn't by itself a good motivation for the Patterned View. If you only cared about agent-neutral good, you'd appropriately aim at *maximizing* the amount of agent-neutral good you produce – you should accept act consequentialism. Act consequentialists could replace the Consilience Requirement with the Act Requirement.

ACT REQUIREMENT: the moral value of \mathcal{R} is the agent-neutrally good or bad of every act that complies with \mathcal{R} .

Motivating the Patterned View requires motivating the Consilience Requirement over the Act Requirement.

CONSILIENCE REQUIREMENT: the moral value of a rule \mathcal{R} is everything actual that's agent-neutrally good or bad to the extent it's caused by actions that \mathcal{R} classifies as morally right.

²¹I take the move from caring about welfare to caring about agent-neutral good more generally to be innocuous; if you disagree, replace my talk about agent-neutral good with talk about welfare.

I'll claim that *Publicity* motivates the Consilience Requirement over the Act Requirement.

Publicity: the two levels of moral thinking are both perspectives on our objective moral reasons.

Given my claim, caring about what's actually agent-neutrally good will combine with *Publicity* to establish the Consilience Requirement.

The Consilience Requirement tells us to ignore certain kinds of agent-neutral good: what's produced by actions that are wrong according to the relevant rule \mathcal{R} . Violating the Requirement thus must involve *paying attention* to the good that the wrong act causes. Remember that we're assuming that there is some set of set of rules that each person accepts at the ordinary level of moral thinking: the rules used when time or interests are limited. Violating the Requirement would thus involve two people: an evaluator (**E**van), and a **p**atient who acts wrongly (**P**at), where:

- at the ordinary level, Evan relies on a rule \mathcal{R} that classify Pat's action as wrong, while
- at the critical level, Evan pays attention to the good produced by Pat's action.

Violating the Consilience Requirement would first involve Evan thinking that Pat's action is wrong, at the ordinary level. That's just what it is for him to accept a rule at the ordinary level that classifies Pat's action as wrong. As a result, it makes sense for Evan to resent and be indignant towards her for what she's doing. The fact that someone acts wrongly is an objective moral reason sufficient for resenting them. So Evan has sufficient objective reasons for resenting Pat.

But if Evan pays attention to the good at the critical level, he would have sufficient moral reasons for being *grateful* for Pat's actions. After all, Pat's actions make the world better than it otherwise would be. And at the critical level, Evan is focusing only on the consequences of behavior – and it makes most sense for a consequentialist like Evan to feel grateful about behaviors that make things go better. So at the critical level, he has objective moral reason sufficient for being grateful to Pat for acting as she does. At this point, Evan has violated *Publicity*: he acknowledges different reasons at the critical level than at the ordinary level.

The violation of *Publicity* is perhaps clearest because the two levels of moral thinking rationalize inconsistent attitudes for Evan. The critical level rationalizes gladness about Pat's behavior, while the ordinary level rationalizes resentment. And those two attitudes are rationally inconsistent. Of course, gladness and resentment about the very same behavior can sometimes be rational. Maybe I can resent an action that's immoral, while also being glad about it because it benefited me – all without irrationality. In this case, though, I'm glad for prudential reasons, while resenting for moral reasons. Even if those attitudes are rational, Evan's attitudes cannot be. Given *Publicity*, the two levels of moral thinking are different perspectives

on Evan's *objective moral* reasons. And it's hard to see how our consequentialist can rationally be glad about some behavior for objective moral reasons while also resenting the behavior for the very same objective moral reasons. For instance, our consequentialist can't appeal to *incommensurability* between the reasons. That appeal would require different reasons at the two levels, which violates Publicity.²²

Evan is entangled in rational inconsistency. And the only assumption that I made about him was that he was violating the Consilience Requirement. So we can conclude that violations of the Consilience Requirement rationalize inconsistent attitudes for someone who accepts Publicity. The only way to avoid rationalizing inconsistent attitudes while accepting Publicity is to accept the Consilience Requirement. The Consilience Requirement guarantees that the two levels of moral thinking are consistent and consilient.²³

Now traditional rule consequentialists often care *both* about something like Publicity and also about what's agent-neutrally good. As Darwall notes, Parfit's shift from the global consequentialism of *Reasons and Persons* to the Triple Theory of *On What Matters* rests on Parfit's coming to accept "a view of the deontic concept of moral rightness that ties it closely to blameworthiness and accountability in a way that effectively concedes a Rawlsian publicity condition" (Darwall 2014, 79). It thus looks like I'm arguing here that Parfit is somehow entangled in rational inconsistency.²⁴

However, my argument doesn't apply to Parfit, because his motivations are subtly different. The argument for Consilience assumes that our ideal moral agent is motivated by what's *actually* agent-neutrally good. And, as we'll see in the next section, Parfit's ideal moral agent cannot be motivated

²²I'm making assumptions here that some might reject; for instance, Richard Arneson (2004) denies that judgments of wrong suffice for negative reactive attitudes. For my purposes, it's enough if Arneson agrees that Evan would have different reasons at the two levels of thinking – that's the crucial point that establishes a violation of Publicity. Such philosophers can see my appeal to the reactive attitudes as a dispensable heuristic.

²³This motivation for the Patterned View dissolves some challenges. An editor objects: "suppose that the only reason numerous acts that are morally right according to R are actually performed is that our society went to great lengths to see it to it that people internalized R. But suppose that the cost in terms of the agent-neutral badness that was incurred by our society in getting people to internalize R far exceeds the agent-neutral goodness that these acts produced. And assume that these acts would not have actually been performed but for the fact that our society went to these lengths." The editor worries that that would mean, implausibly, that R is one of the best rules. The Patterned View would include R as one of the best rules. But it'd also include a rule forbidding the inculcation of R, since the costs outweigh the benefits. The actions that R enjoins are then right, even though inculcating R was wrong. This verdict should look principled given the motivation for the Patterned View: it'd make sense to be grateful to those who *do* the R-actions, since they're part of a good pattern, even though it'd also make sense to resent those who inculcated R.

²⁴Similar questions arise about the proposal from Dale Miller (2020), which also emphasizes something like a Rawlsian publicity condition.

merely by what's *actually* agent-neutrally good; his agent must also care about what's *counterfactually* good. The central motivation for this paper is to formulate pattern-dependent views like rule consequentialism free of counterfactuals, focused only on what's actually good. I take the argument of this section to show that the Patterned View is the *uniquely* correct way to do so. Even if I didn't succeed in establishing that strong of a conclusion, the argument at least suggests that it's a *highly principled* way of doing so.

I've emphasized how schematic the Patterned View is – how it can incorporate the advances that others have made in formulating rule consequentialism. My initial formulation focuses in effect on *following* the rules:

CONSILIENCE REQUIREMENT: the moral value of a rule \mathcal{R} is everything actual that's agent-neutrally good or bad to the extent it's caused by actions that \mathcal{R} classifies as morally right.

But the argument from Publicity suggests a further development of the Patterned View. The reactive attitudes respond to the agent's *quality of will* as displayed in their actions, not just the actions themselves. Given the argument from Publicity, a better Consilience Requirement might restrict attention to actions that *display the agent's quality of will* and that \mathcal{R} classifies as morally right.

4 How the Patterned View is an improvement

We should expect the Patterned View to be an improvement over extant versions of rule consequentialism. It eliminates the counterfactual comparisons that play a central role in those views. And those counterfactual comparisons are just the sorts of things that we know from other subfields create important problems.

This section will show that the Patterned View makes good on this expectation. It focuses on two core problems: the ideal worlds objection that Gideon Rosen (2009), Derek Parfit (2011), and Abelard Podgorski (2018) have recently discussed, and problems about 'minimizing conditions' that David Lyons (1965) introduced and that Holly Smith (2010) generalized. These problems are just the sorts of problems we would expect to arise for views where counterfactuals play a fundamental role. The Patterned View will avoid them precisely because it's free of counterfactuals.

But I want to begin by setting aside one kind of challenge to rule consequentialists. Imagine a demon who'll cause actual disaster unless we all believe that killing our firstborn is required. If the demon's effects happen both in the actual world and in the worlds that matter for a particular rule consequentialist, any kind of rule consequentialist, including the Patterned View, would predict that killing one's firstborn is morally required.

Though this sort of demon example is silly, it elucidates the facts that pattern-dependent views acknowledge as grounds. It shows that pattern-dependent views ground a great many moral facts in contingent empirical facts – many more moral facts than other normative theories. Rule consequentialists can reasonably respond that morality is sensitive to problems

in the real world; morality depends on what would *realistically* make things go best. As a result, moral facts can appropriately be grounded in contingent empirical facts. Fanciful examples about how the actual would could be are not objections, just elucidations of the view. Moreover, *objections* from fanciful examples about what could happen reveals a kind of moral bankruptcy, not taking our practical problems with appropriate seriousness. Rawls emphasizes the same point: “there is no objection to resting the choice of first principles upon the general facts of economics and psychology” (Rawls 1971, 137), and vigorously resists fanciful counterexamples (Rawls 2001, §19). §5 will return to this point.

So I acknowledge at the outset that the Patterned View and traditional rule consequentialists all face a common feature: that they ground a great many moral facts in contingent empirical facts. §§4.1-2 will describe two additional problems for traditional rule consequentialists, and show how the Patterned View solves them. Those two problems are different from and much more serious than any problems about what actually happens. Remember that the ambition of this paper is to introduce the Patterned View as a near-Pareto improvement on extant pattern-dependent views. Even if you think problems about actual effects doom the Patterned View as well as traditional consequentialism, my goal is to convince you that the Patterned View is the most promising formulation.

4.1 The ideal world objection

Parfit introduces the ideal world objection with the pacifist rule: never use violence. Universal acceptance of this rule would plausibly be better than universal acceptance of any other rule about violence, because violence wouldn’t be necessary if no one else is violent. So a universal acceptance consequentialist seems forced to predict that violence is always wrong.²⁵ And that prediction is implausible. The shallow problem is that it’s implausible for those of us who find pacifism facially implausible.

The deeper problem is that the explanation of pacifism seems wrong, that it rests on something irrelevant. In fact, even pacifists should agree that it rests on something irrelevant. For the universal acceptance rule consequentialist, actual violence is simply *irrelevant*: she idealizes in a way that erases actual violence. Defensible versions of pacifism, in contrast, squarely face the violence that actually occurs. A natural diagnosis is that the wrongness of violence must be grounded in *actual* features of violence. Very roughly, the ideal world objection is that universal acceptance rule consequentialism grounds its wrongness elsewhere.

Brad Hooker’s rule consequentialism faces the same kind of problem.

“an act is wrong if and only if it is forbidden by the code of rules whose internalisation by the overwhelming majority of everyone everywhere in each new generation has maximum

²⁵There are lots of tricks it’s tempting to try, but Podgorski works through them systematically and shows that they cannot succeed (Podgorski 2018, 287ff, 293n7).

expected value in terms of well-being (with some priority for the worst off)” (Hooker 2000, 32)

Rawls insists on the fact of reasonable pluralism: that “a plurality of reasonable yet incompatible comprehensive doctrines is the normal result of the exercise of human reason within the framework of the free institutions of a constitutional democratic regime” (Rawls 2005, 16). Hooker’s idealization erases this sort of plurality, or at least insists that no more than an overwhelming minority will accept different moral views.

And the fact of reasonable pluralism plays a central role in Rawls’ reasoning about toleration, a role that would not be played by the facts that Hooker acknowledges about an overwhelming minority. For one thing, it’s much easier for a society to be stable for the right reasons if an overwhelming majority is guaranteed to accept the same moral convictions. Now the details of the problem for Hooker are going to be much more subtle than the details of the problem for the universal acceptance consequentialist, but we don’t need to work through the details to anticipate that similar problems will arise. In both cases, the views erase important features of what we’re evaluating: either the fact of violence or the fact of reasonable pluralism.

It’s natural to hope for a version of rule consequentialism that avoids this sort of problem. We could go through different versions, checking in each case whether the version allows for examples like the pacifism example or the reasonable pluralism example. But that approach risks losing the forest for the trees – missing general patterns in the deluge of examples. A better approach would be to find a general recipe for generating these sorts of examples, and then see whether any version of rule consequentialism can evade the general recipe.

Abelard Podgorski (2018) has recently provided just such a general recipe. He focuses on what he calls *duds*, which are objects that can produce either a great amount of good or a great amount of bad that are only triggered by non-actual patterns of behavior. The pacifism example and the reasonable pluralism example both involve duds. The pacifism example can be modeled as a dud that is triggered only by universal acceptance of the pacifism rule. It’s as if there is an angel or a demon who would make the world extremely good iff everyone complies with the pacifism rule. Since that angel/demon won’t actually do anything, their presence is intuitively morally irrelevant. It doesn’t matter morally whether they make things extremely good or extremely bad iff everyone accepts the same rules, but it matters for universal acceptance rule consequentialism. The reasonable pluralism example involves a dud triggered by internalization by an overwhelming majority of the same rules. It’s as if there is an angel or a demon who would make the world extremely good iff an overwhelming majority internalized the same rules about morality. Since that angel/demon won’t actually do anything, it doesn’t matter morally if they’d make things extremely good or extremely bad, but it would matter for Hooker.

The same problem arises for Parfit’s own view: adopt the rules that do the best at any level of adherence. His view will face the ideal world objection, because it takes duds to be relevant. Assume at first that there

aren't any duds on any level of adherence. Then suppose that a dud is added that makes universal adherence to a rule requiring ϕ -ing very bad. This dud affects what rules do best at any level of adherence, by affecting what does best at universal adherence.

Once we've seen the problem for Hooker and Parfit, we can see that all extant versions of rule consequentialism will mistakenly take duds to be morally relevant. Let \mathcal{P} be the pattern (or set of patterns) that the rule consequentialist takes to determine what's forbidden, required, or permitted. (Maybe \mathcal{P} is internalization by an overwhelming majority in each new generation.) Assume that there is no actual \mathcal{P} of ϕ -ing, either now or in the future, and that there is no dud on a \mathcal{P} of ϕ -ing. (In our example, that'd mean that there are no demons or angels, and that we're not internalizing the exact same rules in each new generation.) Then suppose that a dud on a \mathcal{P} of ϕ -ing is *added*. That dud can flip what was previously required into something wrong, if the dud has bad enough consequences, or flip something previously wrong into something required, if the dud has good enough consequences.

The ideal world objection applies even to views that focus just on the *actual* patterns, like Richard Miller's actual rule utilitarianism:

(ARU2) An act, a , is morally right if and only if it is not the case that there is a practice, P , such that P governs the situation in which a would occur; P forbids a ; P is *de facto* in the society in which a would be performed, and P promotes happiness (Miller 2009, 19).²⁶

Promotion is comparative: promotion as compared to some baseline. There can be duds on whatever baseline is taken to be relevant. Those duds would be triggered at the baseline but could never be triggered in the actual world. Adding those duds would affect what we're required to do on actual rule utilitarianism; the very same problem arises.²⁷ That's why the Patterned View avoids any mention of promotion. Very surprisingly, merely appealing to *actual* effects isn't enough to solve the problem. Podgorski (2018) argues that the ideal-world objection will arise for *any* rule consequentialist view.²⁸

²⁶The view that Conrad Johnson (1991) defends is similar.

²⁷See p. 291 of Podgorski (2018) for his presentation of this point.

²⁸Problems about duds are very different from problems about actual effects on what happens. I earlier imagined a demon who would cause actual disaster unless we all accepted that killing the firstborn was morally required. All kinds of rule consequentialism, including the Patterned View, would take that demon to affect our moral obligations. I suggested earlier that all rule consequentialists should dismiss this case as fanciful: that they insist that morality addresses our actual practical problems, and such demons are not among them. The demon example isn't an objection; it only shows that moral facts are ultimately grounded in contingent empirical facts for rule consequentialists.

Crucially, though, dismissing the actual disaster case as fanciful *requires* a solution to the ideal worlds objection. After all, the ideal worlds objection shows that the targeted view is itself fanciful and ignores the problems in the real world. For example, the pacifism example shows that universal acceptance rule conse-

This paper aims to formulate a pattern-dependent view free of counterfactuals. I said that it's worth looking for such a view because we know that appeal to counterfactuals generates systematic problems. And the ideal world objection is just the sort of problem we would expect to arise – I think it shows that extant rule consequentialists all commit the Conditional Fallacy that Shope (1978) describes. (Someone who understands that fallacy could have generated much of the literature on the ideal world objection just by applying Shope's general recipe.) So when we encounter the ideal world objection, we should think: of course this kind of objection arises – it's just what we would expect to arise given the appeal to counterfactuals! In both literatures, too, silly examples about angels or demons are the easiest ways to illustrate the problems, but the problems are serious even if the easiest illustrations are silly.

The Patterned View cleanly avoids the ideal world objection. The Patterned View is that the demands of morality are the demands of the rules that are morally best, with the:

GENERALITY REQUIREMENT: the morally best rules are sufficiently general, and the
 CONSILIENCE REQUIREMENT: the moral value of a rule \mathcal{R} is everything actual that's agent-neutrally good or bad to the extent it's caused by actions that \mathcal{R} classifies as morally right.

Duds are automatically irrelevant. Actual events are the only things that matter for the Patterned View. Since duds don't change anything actual, duds cannot change what is right or wrong. And unlike Miller's view, *promotion* plays no role, so there is no problem about duds affecting the baseline. No surprises here: the ideal world objection arises essentially from use of counterfactuals, so the Patterned View will of course avoid it.

Now the Patterned View does rest on facts about *causation*, and some philosophers defend counterfactual analyses of causation, as noted earlier. But their appeal to counterfactuals doesn't reintroduce the ideal world objection. Causation is a 'black box' for the Patterned View: as long as your account of causation agrees that fair trials cause stable systems of justice, the Patterned View can incorporate it. Moreover, a counterfactual analysis of causation doesn't itself answer the ideal world objection. It doesn't help Parfit answer it, or anyone else. (If you're skeptical, note that someone who

quentialism fancifully ignores actual violence. Duds may only show that a view *possibly* ignores actual features of an action, rather than showing it *actually* does – but I'll just assume that we don't want even a view that possibly ignores actual features. The rule consequentialist thus faces a dilemma. Her view will have implausible consequences given sufficiently fanciful suppositions about what actually happens. If she acknowledges a need to address those suppositions, we have decisive evidence against her view. But if she dismisses the fanciful suppositions as fanciful, she needs to solve the ideal worlds objection. If she doesn't, her own view allows for similarly fanciful explanations.

§5.2 discusses this kind of point in further detail.

accepts a counterfactual analysis of causation would still think that Parfit has a serious problem.) I earlier claimed that proponents of counterfactual analyses of causation should think that *something special* about causation allows a counterfactual analysis to succeed. I take this subsection to have made good on that claim. The ideal world objection remains an insuperable problem for traditional kinds of rule consequentialism, whatever is true about causation.

4.2 Minimizing conditions

David Lyons (1965) identified a debilitating problem for rule consequentialists – what he calls a problem about ‘minimizing conditions’. Holly Smith generalized the problem, which she introduces with this example.

An important public good would be secured if almost everyone avoided polluting the river. If 90 percent of the industries bordering the river dispose of their waste elsewhere, the river will be healthy. However, if fewer than 90 percent dispose of their waste elsewhere, and more than 10 percent of the industries instead discharge waste into the river, the river will be dangerously polluted. On the other hand, suppose almost all the industries bordering the river discharge their waste into it so that the river is seriously polluted. If you are one of the industrialists, for you to join the few industrialists who avoid discharging waste into the river would be to impose a cost on yourself without producing any public good, since the river will be polluted whatever you do. (Smith 2010, 419)

Since your refusal to discharge into the river makes no difference but is costly for you, discharging into the river seems permissible. But rule consequentialists struggle to explain why. Focus in particular on the cases where some public good exists at the rule consequentialist’s favored embedding but the public good won’t actually exist. Because the favored embedding produces the public good, the rule consequentialist seems to tell you to do your part to producing the public good; in Smith’s pollution case, the rule consequentialist seems to predict that discharging your waste into the river is wrong, rather than permissible. Lyons thinks of this problem as a problem about *minimizing* conditions, because he thinks of it as a problem about minimizing how bad things get for you if others aren’t cooperating.

A natural response is that the best rules are conditional: roughly, don’t discharge if others aren’t, but discharge if enough others are too. But at the favored embedding, the conditionalized rule will be no better than the bare rule enjoining you to not discharge. The public good will exist if either the conditional rule *or* the non-conditional rule is embedded. And since the embedding will be one where others aren’t discharging waste, the part of the conditionalized rule that distinguishes it from the categorical rule is silent. The two rules seem tied; the rule consequentialist seems unable to deliver a determinate answer. Smith shows in compelling detail that it’s

simply impossible to avoid indeterminacy here, no matter how sophisticated the theory becomes (Smith 2010, 421–426).

Now contemporary rule consequentialists often break such ties by preferring whichever rule is closest to conventional morality (Hooker 2000, 32).²⁹ Discharging waste into the river would then be permissible in part because conventional morality permits discharging it. Smith doesn't discuss this response. It's perhaps unsurprising that she doesn't, because it is quite strange. Though it avoids her problem, it's not clear whether it avoids it in the right way. Conventional morality as a tie-breaker makes sense for an *act* consequentialist who sees rules as mere rules of thumb. But it makes much less sense as an essential part of the nature of morality.

In particular, it's hard to see how moral *realists* can use conventional morality as a tie-breaker. It's natural to characterize moral realists as philosophers who take moral properties (like being right or being wrong) as independent of and more fundamental than our evaluative attitudes. Think of a layer metaphor: the lower layers are the more fundamental ones. For moral realists, moral facts are in a lower layer than our evaluative attitudes. But incorporating conventional morality as part of the essence of morality puts our evaluative attitudes in the *same* layer as moral facts. So it abandons moral realism.³⁰

The Patterned View supports a very different answer to Smith and Lyons.³¹ The simplest version of Smith's case is one where we already know that more than 10% of the industrialists are polluting. In that case, the best rule would be:

River, given belief: discharging into the river is right if you reasonably think the public good won't exist.

This rule contrasts with a rule forbidding discharging into the river.

Elsewhere, given belief: discharging into the river is *wrong* even if you reasonably think the public good won't exist.

River, given belief is credited with the benefits of discharging into the river: the benefits that flow from the resources freed up by discharging into the

²⁹Richard Brandt (1963) and David Copp (1995) incorporate similar ideas.

³⁰A referee helpfully suggests that the appeal to conventional morality might serve only to make morality accessible to the agent, so that it's an instance of the sort of perspectivism that T. M. Scanlon (2008) or Elinor Mason (2013) defend. But the appeal to conventional morality has to be doing more than that. We know that conventional morality could have been different – it could have required not discharging into the river even if others discharge. A moral realist would insist that only one of those two possibilities is one where conventional morality gets moral reality right. It seems like Hooker would have to disagree. And he can't trace that disagreement back to a limitation on our current information: we know that conventional morality could have been different.

³¹I'll assume that the Generality Requirement isn't itself enough to answer their objection; they're describing a general possibility by focusing on a specific example, and I'll follow their lead.

river. It wouldn't be credited with the public good. But in the cases we're now discussing, the public good doesn't exist. So the conditional rule gets credited with the only benefits that actually exist, while the *Elsewhere, given belief* rule doesn't. So *River, given belief* is best: the Patterned View predicts that in this case, discharging into the river is the right thing to do. (Take 'reasonably thinks' as a placeholder for a more complete account.)

If the public good will exist, the Patterned View predicts that your discharging into the river is *wrong*, even if your actions make no difference for the public good. The Patterned View is a recipe for evaluating *everyone's* actions, not just your own. You should think that the people who are producing the public good are acting rightly. Since the only difference between you and each one of them is a singular difference – that you're you and they're them – the difference is morally irrelevant: you must think that everyone who performs that action is doing what's right. Lyons and Smith would be more sympathetic with this prediction; they think problems about what Lyons calls *minimizing*-conditions are much more serious than problems about his *maximizing*-conditions.

Turn now to the complicated version of Smith's case. Suppose that sometimes the public good exists and sometimes it doesn't. That is, there are lots of rivers and lots of industrialists, and sometimes the public good exists and sometimes it doesn't. Suppose further that you don't know what sort of case you're in: whether the public good could exist on *your* river.

River, given ignorance: discharging into the river is right if you're reasonably unsure if the public good will exist.

Elsewhere, given ignorance: discharging into the river is wrong if you're reasonably unsure if the public good will exist.

For the Patterned View, the proportions among the cases matter. The public goods that actually exist are credited to *Elsewhere, given ignorance*, while the benefits of discharging into the river are credited to the *River, given ignorance* rule. The public goods are significantly greater than the benefits of discharging. So as long as there are *enough* public goods, the *Elsewhere, given ignorance* rule will be best. One of our considered moral judgments is that we're required to do our part to produce a public good unless we reasonably think it won't exist. The Patterned View vindicates this considered judgment given the further plausible assumption that there are enough public goods for the proportions to work out.

And the Patterned View is fully compatible with moral realism. In the first place, it doesn't put conventional moral judgments on the same fundamental level as moral facts: the only things that are at the same level are facts about what people's actions cause. And any normative theory where effects matter will put them on the same level as the Patterned View. In the second place, the insistence from the last section that morality address our practical problems is compatible with moral realism. No one should think that Philippa Foot (1961) or G.E.M. Anscombe (2000), say, have abandoned realism for insisting on parallel points.

Full discussion of Smith's problem would take another paper. But even before we see all the details, we can be confident that the Patterned View is

an improvement over extant versions of rule consequentialism. Extant versions either predict that we're required to try to produce a public good we know can't exist, or that we must abandon moral realism. The Patterned View avoids both predictions. Interesting further questions shouldn't distract from the progress that Patterned View makes.

5 Wrapping up

I claim that the Patterned View is a near-Pareto improvement over extant rule consequentialist views. I've already shown that it cleanly solves a range of problems. I myself think that it's a *perfect* Pareto improvement; I don't think that it loses any genuine advantages of rule consequentialism. But it does deliver different verdicts about some cases, and you might reasonably think that those differences instead make it a *near* Pareto improvement.

5.1 Act consequentialist rules as subordinated

As background to what I'll say in the following sections, let's review. We're looking for a counterfactual-free view that captures the moral significance of patterns of behavior that together achieve more good than individuals could on their own. And the Patterned View does just that. For instance, it gives a distinctively patterned explanation of the wrong of scapegoating.

But nothing guarantees that in each of our diverse human endeavors, there will *always* be a pattern of behavior that's better than individuals striking out on their own. Maybe in some areas, individuals striking out is what's best – some obligations of beneficence may be correctly explained in that way – so let's assume that individuals striking out is sometimes best. Given that assumption, the Patterned View will also incorporate an act-consequentialist style rule:

- The BENEFIT Rule: choosing a greater benefit over a lesser benefit is right

Crucially, though, the Patterned View will *subordinate* the Benefit Rule to the other rules we've discussed.

- The FAIR TRIALS Rule: scapegoating is wrong even when the odds of discovery are low enough, and fair trials are always right.

It'll subordinate the Benefit Rule by taking it to only apply to actions that already conform to the other rules. So it'll never enjoin scapegoating, because it's subordinate to the rule that forbids scapegoating. The Patterned View might subordinate the Benefit Rule by taking other rules as lexically prior to it, or by reformulating the Fair Trials rule to treat scapegoating as an exclusionary reason that undercuts the significance of its benefits. Think of the range of options for subordinating the Benefit Rule as extra resources for the Patterned View: whichever one works best is justified because it puts our intuitions in reflective equilibrium.

I'm not stipulating that the Patterned View subordinates the Benefit Rule. Instead, the Patterned View *predicts* that the best set of rules subordinate the Benefit Rule to the Fair Trials Rule. To see why it makes that prediction, compare a set of rules that subordinates the Benefit Rule to the Fair Trials Rule to the set that consists only of the unsubordinated Benefit Rule. The Consilience Requirement credits the great good of a stable justice system only to the set containing the Fair Trials Rule. It doesn't credit those benefits to the set just with the *unsubordinated* Benefit Rule. The Consilience Requirement credits a set of rules with the goods caused by actions that that set classifies as right. Our guiding assumption has been that fair trials cause a stable system of justice, but also that individual acts of scapegoating won't themselves destabilize that system. So individual acts of scapegoating themselves are of a greater benefit than individual free trials.³² As a result, the Consilience Requirement does *not* credit the unsubordinated Benefit Rule with the good of a stable justice system, because that good is caused by actions that are *wrong* according to the unsubordinated Benefit Rule. In contrast, the set of rules that subordinates that Benefit Rule to the Fair Trials Rule is credited with the good of a stable justice system *together with* the good of actions that are right according to the *subordinated* Benefit Rule. So that set is best.

5.2 Good potential patterns

The subordinated Benefit Rule answers some important challenges to the Patterned View. A referee helpfully contrasts an extant patronage system for jobs with a merely possible system that awards jobs based on merit. The extant patronage system produces actual good, while the merely possible system doesn't, because it's not in place. But the merit-based system plausibly *would* produce a great deal more good. And the Consilience Requirement credits rules only with actual good – so the referee worries that it won't predict a requirement to switch to the merit-based system. Now 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. So let's suppose that the merit-based system *never* takes hold: nothing but patronage to the end of time.

Crucially, though, the patronage system still violates the Benefit Rule: the patronage system involves individuals choosing a lesser benefit over a greater. So the Patterned View still predicts that choosing based on merit is what's right. But the Patterned View does not give a distinctively patterned explanation of the wrong of the patronage system, since the Benefit Rule doesn't reflect goods distinctive to patterns. The Patterned View takes wide swathes of our considered moral judgments to reflect goods distinctive to patterns, but it doesn't need *every single one* of our considered moral judgments to do so.

³²§1 defended this claim as necessary for the advantages that rule consequentialism purports to secure.

The Patterned View offers a new account of the moral significance of goods distinctive to cooperative patterns. It's the fourth possible account of those benefits. The four possible accounts treat the benefits as:

- irrelevant for what I objectively ought to do (act consequentialists).
- relevant for what I objectively ought to do only if I can help bring them about (Regan (1980); Woodard (2019)).
- relevant for what I objectively ought to do *if the cooperative pattern exists*, even if I can't do anything to bring it about or sustain it (the Patterned View)
- relevant for what I objectively ought to do even if they don't exist and never will (traditional formulations of rule consequentialism or views like Woodard (2008)).

Rule consequentialism promises to put our considered judgments in reflective equilibrium in a way that *act* consequentialism doesn't: it vindicates more of our judgments about justice, promise-keeping, truth-telling... The relevant cooperative patterns *already exist* – so the Patterned View vindicates the judgments that motivate pattern-dependence. Now a core advantage of act consequentialism is its grounding morality in something we can intelligible 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. So given the ultimate motivation for the Patterned View, the referee's suppositions about the patronage system guarantee that obligations to distribute jobs based on merit can't have a pattern-based explanation. The benefits of each act considered on its own are the only thing that could matter, since the pattern never comes to exist.

I've been granting that the patronage system will last until the end of time. But we wouldn't know one way or another; the Patterned View needs fleshing out to guide ignorant creatures like us. Some versions of the Patterned View could ground our moral obligations in patterns with a reasonable chance of being actual. Such versions would classify patronage as wrong so long as there's a reasonable chance that a merit-based alternative comes to exist. Those versions go even further in eliminating status quo bias. (In any case, reasonable chances of becoming actual must determine what's right and wrong if the future is indeterminate.) The Patterned View also needs fleshing out since the best rules could tie. For one thing, the rule 'always act rightly' would tie with the best *specific* set of rules forbidding scapegoating and the like. I'd flesh the View out by integrating it with a metaethical view I've defended elsewhere (2017), which in effect lowers the bar for what it takes to see a set of rules as best.

Fully evaluating status quo bias in the Patterned View would thus take another paper. We'd need to investigate the plausible ways of developing

the View. As emphasized, the View rebuilds rule consequentialism from the ground up. I've tried to display the basic advantages of the View in enough detail to make us hopeful about its prospects. But some versions of the View will have further advantages.

5.3 On Hobbesian worlds

The Patterned View promises to put our considered judgments in reflective equilibrium because patterns of fair trials, promise-keeping, truth-telling, and the rest already exist. It's thus reasonable to wonder about worlds where those patterns never took hold. Imagine a Hobbesian world of interminable war of all against all – a world where cooperative patterns never have existed and never will. The Patterned View collapses into a kind of act consequentialism at that world. There aren't any actual patterns of group behavior that stand as grounds to distinguish the Patterned View from a kind of act consequentialism. Similarly, a world where only two sentient beings exist but never interact would be a world where something like the Benefit Rule alone ends up best, with the two beings then having act-consequentialist-style obligations.

Traditional rule consequentialism might work differently: it might imagine how things would go at a non-Hobbesian world. Hobbesian worlds make me concede that you might see the Patterned View as merely a *near*-Pareto improvement. But two points are important in response. The first point is that, within the Hobbesian world, attention to the non-Hobbesian world is just like attention to worlds where everyone is a pacifist: in both cases, there's no path from here to there. Trying to conform your behavior to the non-Hobbesian world or the pacifist world seems fetishistic. But if you insist on a fetishistic moral theory, you would see the Patterned View as merely a *near* Pareto improvement, rather than a perfect Pareto improvement.

The second point is that intuitions about the Hobbesian world cannot be genuine evidence against the Patterned View. All rule consequentialists should insist that intuitions about the Hobbesian world are nothing more than *projections* of our intuitions about the actual world onto the Hobbesian world. Supposing otherwise would involve supposing that our intuitions arise partially from an ability to track the effects of different embeddings of rules. And I've argued elsewhere (2019) that rule consequentialism is self-defeating if it rests on that supposition.³³ My (Perl) focuses specifically on intuitions about Hobbesian worlds, arguing that they won't provide evidence one way or another; they're just projections. Since those intuitions are just projections, the Patterned View explains them the same way that everyone should.

5.4 Looking forward

This paper has provided the first explicit formulation of pattern-dependence free of counterfactuals. I emphasized at the outset that appeal to counter-

³³My (2017) shows how rule consequentialism can be developed so it doesn't rest on that supposition.

factuals is problematic in itself – that lots of highly general problems arise simply from that appeal. But I also promised that more familiar problems would have clean solutions – and I think the discussion of minimizing conditions and of the ideal world objection made good on that promise, with a range of other further interesting applications saved for another occasion.

The Patterned View also matters as a model for reformulating other pattern-dependent views, like Scanlon’s contractualism. The Patterned View replaces counterfactual evaluation of rules and principles with *reflective* evaluation: credit rules and principles with features of actions that are right according to those very rules and principles. Rule consequentialism is one of the simplest pattern-dependent views, where agent-neutral value is the only thing that matters. So if we can formulate the simplest pattern-dependent view freed of counterfactuals, we can reasonably expect to learn lessons that help us formulate more complicated alternatives too, because we’re coming to better understand the range of ways to formulate pattern-dependence.

References

- Anscombe, G. E. M. (2000). *Intention*. Cambridge, MA: Harvard University Press,.
- Arneson, R. (2004). Moral Limits on the Demands of Beneficence? In D. Chatterjee (Ed.), *The Ethics of Assistance: Morality and the Distant Needy*, pp. 33–58. Cambridge University Press.
- Berkeley, G. (1712). Passive Obedience, or the Christian Doctrine of Not Resisting the Supreme Power, Proved and Vindicated upon the Principles of the Law of Nature. In D. Monro (Ed.), *A Guide to the British Moralists*, pp. 217–27. London: Fontana.
- Berker, S. (2018). The Unity of Grounding. *Mind* 127, 729–777.
- Berker, S. (2019). The Explanatory Ambitions of Moral Principles. *Nous* 53 (4), 904–936.
- Brandt, R. (1963). Toward a Credible Form of Utilitarianism. In H.-N. Castañeda and G. Nakhnikian (Eds.), *Morality and the Language of Conduct*, pp. 107–43. Detroit: Wayne State University.
- Brandt, R. (1967). Some Merits of One Form of Rule-Utilitarianism. In *University of Colorado Studies in Philosophy*, pp. 1–22. University of Colorado.
- Copp, D. (1995). *Morality, Normativity, and Society*. New York: Oxford University Press.
- Darwall, S. (2014). Agreement Matters: Critical Notice of Derek Parfit, On What Matters. *The Philosophical Review* 123 (1), 79–105.

- Dreier, J. (1993). Structures of Normative Theories. *The Monist* 76 (1), 22–40.
- Feldman, F. (1975). World Utilitarianism. In K. Lehrer (Ed.), *Analysis and Metaphysics*, pp. 255–273. D. Reidel Publishing Company.
- Fine, K. (1994). Essence and Modality. In J. Tomberlin (Ed.), *Philosophical Perspectives 8: Logic and Language*, pp. 1–16. Atascadero: Ridgeview.
- Fine, K. (2001). The Question of Realism. *Philosopher's Imprint* 1, 1–30.
- Foot, P. (1961). Goodness and Choice. *Aristotelian Society Supplementary Volumes* 35(1), 45–80.
- Hare, R. M. (1952). *The Language of Morals*. Oxford: Oxford University Press.
- Hare, R. M. (1981). *Moral Thinking: Its Levels, Method, and Point*. Oxford: Clarendon Press.
- Harsanyi, J. (1977). Rule Utilitarianism and Decision Theory. *Erkenntnis* 11, 25–53.
- Hooker, B. (2000). *Ideal Code, Real World: A Rule-consequentialist Theory of Morality*. Oxford University Press.
- Hooker, B. (2016). Wrongness, evolutionary debunking, public rules. *Etica E Politica* 18 (1), 135–149.
- Johnson, C. (1991). *Moral Legislation*. Cambridge University Press.
- LazariRadek, K. D. and P. Singer (2010). Secrecy in consequentialism: a defense of esoteric morality. *Ratio* 23 (1), 34–58.
- Lyons, D. (1965). *Forms and Limits of Utilitarianism*. Oxford: Oxford University Press.
- Martin, C. (1994). Dispositions and Conditionals. *Philosophical Quarterly* 44, 1–8.
- Mason, E. (2004). Consequentialism and the Principle of Indifference. *Utilitas* 16 (3), 316–321.
- Mason, E. (2013). Objectivism and Prospectivism about Rightness. *Journal of Ethics and Social Philosophy* 2.
- Miller, D. (2020). Moral Education and Rule Consequentialism. *The Philosophical Quarterly* 71 (1), 120–40.
- Miller, R. (2009). Actual Rule Utilitarianism. *Journal of Philosophy* 106 (1), 5–28.

- Moore, G. E. (1903). *Principia Ethica*. Cambridge: Cambridge University Press.
- Norcross, A. (1990). Consequentialism and the Unforeseeable Future. *Analysis* 50 (4), 253–256.
- Parfit, D. (1984). *Reasons and Persons*. Oxford: Clarendon.
- Parfit, D. (2011). *On What Matters*, Volume 1. Oxford: Oxford University Press.
- Pearl, J. (2000). *Causality*. Cambridge: Cambridge University Press.
- Perl, C. Some question-begging objections to rule consequentialism.
- Perl, C. (2017). *Positivist Realism*. Ph. D. thesis, University of Southern California.
- Perl, C. (2019). Empirical ignorance as defeating moral intuitions? A puzzle for rule consequentialists (and others). *Analysis* 79 (1), 62–72.
- Podgorski, A. (2018). Wouldn't it be Nice? Moral Rules and Distant Worlds. *Nous* 52 (2), 279–294.
- Portmore, D. (2011). *Commonsense Consequentialism*. Oxford University Press.
- Railton, P. (1984). Alienation, Consequentialism, and the Demands of Morality. *Philosophy and Public Affairs* 13 (2), 134–171.
- Rawls, J. (1971). *A Theory of Justice*. Cambridge, MA: Harvard University Press.
- Rawls, J. (2001). *Justice as Fairness*. Cambridge, MA: Belknap Press.
- Rawls, J. (2005). *Political Liberalism*. New York: Columbia University Press.
- Regan, D. (1980). *Utilitarianism and Co-operation*. Oxford: Clarendon Press.
- Ridge, M. (2006). Introducing Variable-Rate Rule-Utilitarianism. *Philosophical Quarterly* 56, 242–53.
- Rosen, G. (2009). Might Kantian Contractualism be the Supreme Principle of Morality? *Ratio* 22, 78–97.
- Rosen, G. (2010). Metaphysical Dependence: Grounding and Reduction. In B. Hale and A. Hoffman (Eds.), *Modality*, pp. 109–136. Oxford: Oxford University Press.
- Scanlon, T. M. (1982). Contractualism and Utilitarianism. In S. Cahn and J. Haber (Eds.), *20th Century Ethical Theory*, pp. 647–665. Prentice Hall.

- Scanlon, T. M. (1998). *What We Owe to Each Other*. Cambridge, MA: Harvard University Press.
- Scanlon, T. M. (2008). *Moral Dimensions: Permissibility, Meaning, Blame*. Cambridge, MA: Harvard University Press.
- Schaffer, J. (2009). On What Grounds What. In D. Chalmers, D. Manley, and R. Wasserman (Eds.), *Metametaphysics*. Oxford: Oxford University Press.
- Shope, R. (1978). The Conditional Fallacy in Contemporary Philosophy. *Journal of Philosophy* 75, 397–413.
- Singer, P. (1972). Famine, Affluence, and Morality. *Philosophy and Public Affairs* 1 (3), 229–243.
- Singer, P. and K. D. LazariRadek (2016). Doing our best for hedonistic utilitarianism. *Etica E Politica* 16 (1), 187–207.
- Smart, J. J. C. (1973). *Utilitarianism: For and Against*. New York: Cambridge University Press.
- Smith, H. (2010). Measuring the Consequences of Rules. *Utilitas* 22, 413–33.
- Thomson, J. J. (1996). *Rights, Restitution, and Risk: Essays in Moral Theory*. Cambridge, MA: Harvard University Press.
- Woodard, C. (2008). *Reasons, Patterns, and Cooperation*. Routledge.
- Woodard, C. (2019). *Taking Utilitarianism Seriously*. Oxford University Press.
- Woodward, J. (2003). *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford University Press.