

Partial Compliance Theory

Henrik D. Kugelberg & Zofia Stemplowska

*Forthcoming in David Copp, Connie Rosati and Tina Rulli (eds.).
The Oxford Handbook of Normative Ethics. Oxford University Press.*

People often fail to comply with the demands of morality. Partial compliance theory takes this noncompliance or its possibility into account in the formulation of moral requirements for people to comply with, or in the evaluation of people's actions against those requirements. This chapter critically engages with recent work on partial compliance theory. It examines the relationship between noncompliance and injustice, assesses different ways of doing partial compliance theory, sketches the relationship between partial compliance theory and other central methodological debates in moral and political philosophy, and considers whether it is reasonable to issue moral requirements without examining what would happen if everyone did their part.

Keywords: Injustice; methods in moral philosophy; non-ideal theory; noncompliance.

1. Introduction

Morality sets standards for people and people fall short of them. Thankfully, the really important standards elicit widespread compliance. Most people do not murder, and most states do not launch wars of aggression. But morality also requires people not to take too long to reply to emails, not to lose due patience with family members and not to waste water. Morality also does not license people to consider only themselves or their family when their decisions have significant implications for others. Our world is not characterised by full compliance with the requirements of morality.

This chapter concerns what moral theory should make of this fact. It discusses moral *partial compliance theory*.¹ Partial compliance theory takes noncompliance, or its possibility, into account in formulating moral requirements or in evaluating people's actions against those requirements.²

Before we begin, a few clarifications are in order. Noncompliance comes in various forms. This poses different problems for partial compliance theory. Noncompliance may be more or less widespread, depending on how many people fail to comply, may be more or less severe depending on how important are the requirements that are violated, and may have different sources depending on why it arises. The term 'noncompliance' itself may reference only wrongful conduct or any conduct, however innocent, that falls short of the prescribed standard. Unless we specify otherwise below, we assume that the noncomplying conduct is wrongful conduct.

Partial compliance theory may be both about trying to determine what substantive requirements apply to people, given that other people fail to comply with the requirements, and about offering systematic reflection on how to take people's noncompliance into account when building substantive theories that issue

¹ We use the term theory loosely, to refer to any attempt to organise some statements (inputs into the theory) in a coherent way in order to arrive at some further statements (outputs of the theory).

² Note that other kinds of theories could also be partial compliance theories. We can imagine a partial compliance theory about the requirements of rationality or prudence or efficiency, and so on. Our focus here is limited to moral theory and requirements of morality.

requirements. In practice, however, those who offer good substantive theories often also offer systematic insights into how to take noncompliance into account, and those who aim to offer such systematic insights tend to offer them as part of substantive theories.³ In what follows, we will reference both types of work, as well as work that combines both perspectives.

Cutting across all of these distinctions, noncompliance may also enter the theory at two different points.

First, a theory might ask: what should those who plan to comply do, given the noncompliance of others? Just war theory, for example, issues requirements for an attacked state on the assumption that it will comply with them, even if the attacker will not. *'L'enfer c'est les autres'*, as Sartre puts it.

But a theory may also be constructed to issue requirements on the assumption that those to whom they are issued will be prone to compliance failures, which the theory hopes to predict and take into account. The question here is not so much about what some should do, given that others will not do their part. Rather, it is about how to design requirements given the fact that people will often – or regularly – fail to do their part. For example, some of the discussion of early anti-Covid regulations during the pandemic lockdown concerned what people could be expected to do, given the inclination to just continue doing what they have always done. Public choice theories – theories developed by economists to deal with problems of governance and public policy – also often assume that public officials will inevitably be tempted to act in their self-interest rather than in the spirit of their mandates. These theories aim to identify the incentives that can be offered to them to get the right results.

Of course, a theory can also attempt to answer both questions. For example, a theory can assume that no one will comply perfectly but ask what requirements it makes sense to issue to those who are more willing to comply than to those who are less willing to comply.

Finally, the term 'moral requirements' can be understood more or less broadly. In what follows we will sometimes also refer to requirements or duties of justice. Different theorists have different views on how morality relates to justice. We adopt a simple view here. When we talk about requirements (or duties) of justice, we mean the sub-set of moral requirements that regulate what persons are owed, rather than what it would be nice or good for them to have. Often what they are owed is something that could be delivered only through the coordinated or mutually-responsive action of multiple individuals.

The aim of this chapter is to offer a systematic overview of how moral theory has attempted to deal with the fact that people tend to fail to comply with the requirements of morality. We examine the relationship between noncompliance and injustice, assess different ways of doing partial compliance theory, sketch the relationship between partial compliance theory and other central methodological debates in moral and political philosophy, and consider whether it is reasonable to issue moral requirements without examining what would happen if everyone did their part.

We begin with the question of whether an injustice or moral wrong exists only when there has been noncompliance. We then consider partial compliance theories that aim

³ There is also a related genre – that of a partial compliance manifesto: these are texts that beseech those who theorise to do it in a particular compliance mode but are not themselves substantive theories or offer systematic insights on how to take noncompliance into account in such theories.

to issue requirements for compliers, given others' noncompliance, before turning to partial compliance theories that expect noncompliance from all. Finally, we explore the more general question of whether partial compliance theory can be done in separation from full compliance theory.

2. Noncompliance and injustice

Can there be injustice without failures of compliance? Some use the term 'structural injustice'⁴ to refer to injustice that arises through a myriad of complex interconnections and interlocking actions, making it difficult to determine who failed to comply with which specific requirements. Many failures of policy in democratic states are of this type: plenty of people act wrongly in letting rivers become contaminated with sewage, but it is difficult to pinpoint wrongdoing with precision, at least without a lot of effort. More generally, as feminists and racial justice theorists have long pointed out, patriarchy and supremacy operate in ways that make separate instances of noncompliance – due to misogyny, racism, or wilful ignorance, or to less pernicious causes – amplify one another to deliver widespread oppression (Nutti 2019, Táíwò 2022).

Understood in this first way, structural injustice involves noncompliance and individual wrongdoing. A key feature is the difficulty of assessing who failed to comply with which requirements.

On a second way of defining 'structural injustice,' however, it is injustice that exists despite the complete absence of individual wrongdoing and noncompliance. Can such injustice exist?

Think of Iris Marion Young's discussion of the single mother Sandy who through various processes 'faces the prospect of homelessness' (2011, 43–4; see also McKeown 2021). A developer buys the property she lives in, and she does not have many good options for other places to live. She finds an apartment far from her place of work, and so she must buy a car for the commute. However, after the car purchase, she cannot afford the deposit for the apartment. She realises that a three month deposit is a 'typical landlord policy,' and so she now risks homelessness.

The people close to Sandy all act in morally permissible ways. In Young's telling, many people even go out of their way to help Sandy.

Does this mean Sandy's situation is an injustice without noncompliance? Not necessarily. Even if there is no noncompliance from the landlords, the developer, and so on, the noncompliance might be located higher up. Perhaps the local government should have earmarked more resources for people in situations like Sandy's. Or perhaps it failed in its duties to provide affordable housing or an adequate and efficient public transport network. Perhaps the noncompliance is located even higher up. That is, we might think high-level public officials failed to bring about a just economic system that ensures that people are not as economically vulnerable as Sandy is.

But suppose Sandy's situation genuinely could not have been predicted or averted. We can imagine a variation of the case, where she finds herself close to homelessness because of a chance event no one could have reasonably foreseen or anticipated. In this

⁴ The seminal recent discussion of structural injustice is in Young 2011. For more on structural injustice see Reiman 2021, Barry and Macdonald 2016, Parekh 2017, Abdel-Nour 2018, Sangiovanni 2018, McKeown 2021, Goodin and Barry 2021, Butt 2021.

case, we may either accept that this is an injustice without noncompliance. Or we might think that there initially is no injustice at all, precisely because no one could have done anything to remedy the situation. Note, however, that if people do not do anything to assist Sandy, once she is in this regrettable situation, that may itself be an injustice. This can be explained by the noncompliance with moral requirements to help those in need.

The general difficulty of identifying cases of injustice without noncompliance is that even if we think that complex systems can generate results that no participants intended, agreed with, or anticipated, we then face the further question whether to describe such results as injustice when they arise or only once they fail to be fixed.

Now, consider a case of someone who is innocent but imprisoned (Overland and Barry 2011). Suppose that the justice system has abundant resources and its officials have acted impeccably. Suppose also that morality requires imprisoning dangerous people.⁵ Suppose further that no institutions could fulfil their roles if they could act only when there was zero chance of making a mistake. Even on these assumptions, an innocent person could be imprisoned. Perhaps, then, this is an injustice without noncompliance.

Some people might want to resist the idea that an innocent person can be convicted without noncompliance. They could insist that when innocents are convicted, there must have been noncompliance somewhere. Following Derek Parfit's distinction (2011:150-62), they might suggest that the officials who imprison the innocent person violate their fact-relative duties even if they act correctly on their evidence-relative duties. Whether or not this strategy succeeds depends on whether it is plausible to say that an action was subjectively right, but objectively wrong.

Is there an injustice here in some other sense? One view is that the term 'injustice' should be reserved for cases where there is a complaint against noncompliance. But the prisoner in our example arguably has no complaint against those who imprisoned her, given that they acted in accordance with their obligations under the criminal justice system.⁶

One might observe, however, that despite the absence of noncompliance, the prisoner has been imprisoned as a result of others' deliberate actions, rather than as a result of non-agential forces. The case of the innocent prisoner thus differs morally from the case of someone who has been trapped in a cave by a falling boulder. After all, unlike falling boulders, fully compliant agents whose deliberate actions adversely affect others may have remedial duties. For this reason, an alternative view might maintain that the term 'injustice' should include not only cases of wrongdoing but also at least some cases where deliberate agency brings about terrible outcomes for others. Still, in the vast majority of cases, and perhaps in all cases, injustice occurs through actions that do not comply with moral requirements. We turn next to partial compliance theory's attempts to reckon with this fact.

⁵ This is of course not a defence of any actually existing prison system.

⁶ She may have it against the wrongdoers that give rise to the existence of a prison system. In related debates some theorists have argued that injustice can arise due to the violation of joint oughts or what Estlund (2019) calls plural requirements that can be violated even if no individual have violated any individual requirements (cf. Dietz 2016, Schwenkenbecher 2021).

3. When others fail to comply

Consider Peter Singer's case in which ten children are drowning in a pond and ten adults are on the shore. Each adult could rescue all the children (Singer 1972, 2002). Suppose that, from the moral point of view, each adult ought to rescue one child. If nine adults leave without rescuing, what is required of the remaining adult? By exploring questions such as this one, theorists try to get a clearer grasp on what our requirements are when others fail to comply, especially by comparing how the answers might vary, depending on, for example, who is in dire need and why, who fails to respond and why, or how frequent and predictable such emergencies are (Unger 1996, Kamm 2000, Arneson 2004, Cullity 2006, Owen 2016).

Sometimes, the noncompliance of others will entail that the compliance of some will not make a difference for the outcome. We could picture a version of the pond case, for instance, where rescuing even a single child requires the coordinated effort of all ten adults. This is not the kind of case we will be discussing. Rather, in what follows, we will focus on cases where compliance by some in the face of noncompliance by others can make a difference to the outcome.

3.1. Nothing special about noncompliance?

Some suggest that noncompliance doesn't alter the obligations of compliers. In Singer's ten-person pond case, for example, it makes no difference to whether the one remaining adult needs to rescue all ten children because the other nine adults left or because a tree fell down blocking their access to the shore. Of course, problems arising due to noncompliance could affect what means would be available for rescue. But noncompliance is not special over and above that, according to what we will call the "makes-no-difference view", or MND.

MND seems to give the intuitively correct answer in the pond scenario: the remaining adult should rescue all ten children whether noncompliance or a fallen tree leaves only her in a position to do so. But can this intuition be vindicated?

That depends on our other theoretical commitments.

Act-consequentialists believe that people are required to do what would bring about the best consequences. On a straightforward reading of what this would entail in the pond case, the best outcome is that all ten children are rescued. When all ten adults are at the shore, the best means for achieving this is that every adult saves one child each. If nine adults leave, whether the complier alone can rescue every child will depend on various factors. Assuming that she can, then saving all ten children is what she must do. In this case, the fact that nine leave does in one sense not change the best outcome, just the best means.

There is a complication, however. Even though act-consequentialists will insist that a person must act to produce the best outcome, what counts as the best outcome will partly depend on the cost of different courses of action. When an adult must rescue the children alone, this can affect the costs she must bear. The effort required of a single rescuer of ten children can outweigh the aggregate effort required of ten adults who rescue one child each. If every adult did their part, they could therefore jointly achieve a better – less costly – outcome.

For an act-consequentialist, when others fail to comply, this noncompliance figures in the calculus just as any other element of the picture: we simply ask how it affects

what outcomes are now possible (taking their costs into account) and then decide which one is the best. In that sense, there is nothing special about noncompliance, and this is the sense in which it makes no difference. It makes no difference to the moral calculation whether people either are unable to or consciously choose not to comply.⁷ But the fact that others do not help out can make a difference – by increasing the costs that compliers must absorb; noncompliance can entail that what previously was the best goal is suboptimal or impossible to reach.

MND is also available to deontologists. Deontologists do not think that the rightness of actions depends only on the best outcome; other things matter too.⁸ For example, deontologists standardly believe that people's moral status gives them some permissions not to act to bring about the best outcome. They may therefore not have to absorb high costs to prevent harm to another. But given reasonable assumptions, deontologists could argue that the costs to the remaining adult at the pond do not outweigh the good of the saved lives. So long as costs to the complier remain reasonable, they have a duty to rescue, regardless of whether the noncompliers don't comply by choice or due to intervening events.

3.2 Noncompliance and unfairness

Contrary to MND, proponents of what we will call the "fair share" view (FSV) contend that morality requires only that compliers do their fair share. Noncompliance may require them to act in new ways but not to absorb more than their fair share of the costs.⁹

FSV comes in various forms and has various justifications. Versions have been developed by L. Jonathan Cohen (1981), Liam Murphy (2000) and David Miller (2011, 2016). Miller (2016) has argued, for example, that seeing noncompliance as special allows us to recognise that the noncompliers, unlike trees, are moral agents: they, rather than the compliers, bear the moral responsibility for their moral failures. It would therefore be wrong to impose the additional costs on compliers. In response to this, it can be argued that people do not standardly think that because wrongdoers must bear responsibility for their actions, we should not aid those they treat wrongly. Contra Miller, we may therefore think that the moral status of noncompliers isn't really called into question when we recognise that their actions affect what compliers must do, including when it means they must bear greater costs.

We may worry that if noncompliance by others simply translates into more costly requirements for compliers, then morality excessively burdens compliers. Although noncompliers are not morally off the hook, by requiring compliers to step up in order to limit the bad the consequences of the noncompliers' failures to act, morality does to some extent make matters better for noncompliers and worse for compliers.

Murphy (2000) has offered the most sustained consequentialist defence of FSV (though it is a view deontologists could adopt as well). According to Murphy, duties of

⁷ This simplifies things somewhat. We could of course imagine that the long term consequences would be better if we discouraged noncompliance in general. We leave this aspect aside.

⁸ A distinction can also be drawn between deontologists and absolutists who do not believe that outcomes matter at all to the question of what morality requires but I put this aside since it is not a remotely plausible view if generalised.

⁹ A variation on the Fair Shares view (identified and criticized in Tadros 2016) is that noncompliance may force people to bear more than their fair share of costs but not as much as they would be required to bear were the problem due to something other than noncompliance.

beneficence are collective: they are not imposed on individuals as individuals but on the collective made up of anyone who is in a position to assist. Such duties can be contrasted with duties not to harm. A person should not be (wrongfully) harmed by anyone, and each individual is under a duty not to harm another person. But each person does not have individual duties to rescue someone in need of it. Rather, the duty to rescue applies only to those who are in position to help. And what each is required to do is whatever is the fair share of what the collective is required to do. Therefore, noncompliance by some with duties of beneficence does not affect what compliers are required to do.

One natural response to Murphy is to grant that duties of rescue are to be derived fairly from the collective duties, in the way he suggests, while insisting that we need to re-do this each time someone must be rescued, and each time someone fails to comply.¹⁰ On this version of the view, when some fail to comply, it changes the content of the individual duties of others when the derivation is re-run. In the pond case, for instance, the collective duty remains to rescue ten children.¹¹ But if five adults leave, we derive new individual duties for the remaining five, so that each now has an individual duty to save two children each.

Murphy would reject this modified version. His view entails that we must pick a particular point in time – call it T_1 – and examine what individual duties can be derived at that moment. If T_1 is when ten adults are standing around the pond, the fact that five adults leave at a later point does not change the duties of the remaining five adults. They retain a duty to save one child each. From T_1 onwards, we cannot update the content of the individual duties.

The first problem with this is that it is impossible to find any non-arbitrary way of selecting T_1 . Suppose one adult arrives to the pond at 10.05, another twenty seconds later, and so on. At what point should we start deriving duties?

The second problem is that it leads to counterintuitive conclusions. If the individual duty is a once and for all fixed derivation of the collective duty, then the derivative individual duty should remain fixed when the need for assistance diminishes. Suppose that while the ten adults stand by the pond one of them unilaterally saves all ten children at no detectable cost to herself. If, as Murphy holds, what one is under a duty to do does not depend on what others do, the remaining nine people seem to remain under a duty to save a child each – even though there are no children left to save.¹² If instead we adopt the more natural view that the individual duties must disappear because it is no longer possible to fulfil them, then we already accept that the duties that fall on individuals are defined not simply with reference to one's initial fair share but also with reference to what remains to be done in light of the actions of others (Stemplowska, 2016: p. 595).

¹⁰ The next two paragraphs draw on Stemplowska 2016.

¹¹ In fact, neither Murphy nor Miller argue that there are no moral requirements for the adults to step in when it comes to pond cases. Murphy thinks that in some special emergency rescue cases his principle may not apply but it would apply to repeated emergencies for example. Miller argues that while there are no duties of justice in such cases there are humanitarian obligations. Fairness can be invoked to do other work than what we discuss here. Nevertheless, we will continue using this stylized case for clarity. Hoesch (2018) has argued that while there is a duty to do more than one's fair share when others are sufficiently close to us, we also have permission to deter them from getting close (but see Stemplowska 2019). See also Karnein, 2014.

¹² We owe this point to Victor Tadros. See his Tadros, 2016.

Murphy attempts to render noncompliance special by making the desideratum of fairness between all those to whom the duties apply an integral factor determining what each individual is required to do. But it is difficult to see why fairness would play this role. Moral requirements should be fair, but it is not easy to defend as a general principle that morality must make fairness a constraint on what duties of beneficence require of people. It would be better, of course, if requirements to rescue those in dire need were fairly divided and fully satisfied, but unfairness may be the price that has to be paid under conditions of partial compliance.¹³

Even where considerations of fairness do not affect the status of noncompliance, noncompliance could still be significant. Some theorists maintain that it is worse when a harmful outcome is a result of someone's wrongdoing than when it arises otherwise. So it is worse when it results from noncompliance than when it does not (e.g., Butt 2019; 2021). Someone's becoming homeless is worse when due to arson than to being struck by lightning. But the alleged fact that the noncompliant outcome is worse doesn't obviously affect the normative requirements for compliers. Suppose that five people end up in hospital because someone used a sick dog to infect them with rabies, whereas five others merely stumbled upon an infected dog.¹⁴ Even if the former five were wronged and therefore treated worse than the other five, surely each of the ten individuals deserves an equal chance of timely treatment. At the same time, the fact of noncompliance might give rise to additional duties and requirements. A particularly clear example is that those who fail to comply may be under a duty to compensate those who were harmed by their noncompliance.

4. When no one can be assumed to comply

Let's turn now from partial compliance theories that aim to issue requirements for compliers to those that assume that all of those to whom its requirements apply may fail to comply.

A first difficulty concerns what noncompliers are permitted to do once they have failed to comply, and what costs they must absorb given their noncompliance. Suppose, for example, that an international organisation with the power to coerce states to admit refugees is made up of noncompliers (Owen 2018). Can the noncompliers permissibly coerce other (noncompliers) even though their own actions help create the need for coercion? We may think that noncompliers may not coerce other states into admitting the refugees they themselves have refused to admit. Or we may think those other states can permissibly be coerced, given that they do have a duty to admit refugees. The difficulty lies with the idea that noncompliers wrong those they coerce, if they coerce instead of admitting refugees themselves, but that the coerced states nonetheless cannot permissibly resist coercion, given that they, too, have duties to the refugees

¹³ Still, there may be something special about noncompliance, and fairness might in some cases limit what compliers must do (Sameer Bajaj 2019). That said, for deontological duties of aid to be triggered it is not enough that some action could deliver a better (in fact best) outcome. Agents standardly have a generous leeway to decide for themselves what ends to pursue. This means that they come under duty to aid others only when the costs to them are sufficiently low and the interests of others are very weighty. If the interests are very weighty and the costs sufficiently low than the fact of unfairness is unlikely to block the duty to serve those interests.

¹⁴ We are grateful to Victor Tadros for this example.

(Stemplowska 2019). A theory addressed to the noncompliers must resolve difficulties with this structure.

Another difficulty concerns what noncompliance to anticipate. If we assume that people will never act any better than they currently do, have we not given up too much normative territory already? This difficulty is faced by one strand of “realist theory.” This strand aims to issue requirements for people as they are, and so it needs to explain what level of partial compliance it is not arbitrary to assume.¹⁵ Partial compliance theory needs to offer a coherent and plausible account of what level of compliance it plans to work with and why. This is not an easy task.

What does it mean to ‘assume’ levels of compliance and noncompliance in the first place? Substantive theories tend to consider what level of compliance is appropriate rather than simply assume some level. Consider John Rawls’s theory of justice (1999a). Political philosophers commonly describe Rawls’s theory as ‘assuming’ full compliance, but this is an inaccurate description of his theory taken as a whole. Rawls does not treat full compliance as a given. Rather, he asks under what conditions full compliance could be sustained, given what we know of people’s capacities. It would not make sense to pose this question if he treated full compliance as a given. Rawls follows Jean-Jacques Rousseau (1762/2002: 149) who famously proclaimed that he will take men ‘as they are’ and laws ‘as they can be made to be’.¹⁶

Partial compliance theory of this type also needs to be alert to whose perspective it adopts when issuing requirements and what it can justify from this perspective. Rawls’s theory deems inequalities of income and wealth to be permissible when they are necessary for ensuring the most advantage for the worst off. This is interpreted as permission for the market to offer incentive payments to highly skilled workers who demand higher salaries, if and when offering such incentive payments is necessary to get the workers to work in the service of improving the position of the worst off. In response to this, G.A. Cohen (2008: 38–41; see also 1997) has argued that if the highly skilled workers do not need the higher salaries to take the jobs, and if they can be criticised for demanding higher salaries, then Rawls’s theory does not offer a justification for resulting inequality: the inequality is necessary from the employers or policy makers point of view, but it is not necessary from the highly skilled workers point of view. In some ways, they act like kidnappers demanding a ransom. One could, of course, disagree with Cohen’s criticism of Rawls, while accepting his advice to be careful regarding which perspective a theory adopts and what it actually ends up justifying.

The discussion so far has assumed that failures of compliance should be classified as wrongdoing. But which actions count as (wrongful) noncompliance is a difficult question. It can be difficult to know the limits of people’s capacities, as well as how those limits bears on what can be classified as wrongful. If the average person cannot run the 4 minute mile or multiply quadruple digits in their mind, they clearly could not be morally required to do so. In the vocabulary developed by Onora O’Neill (1988, 1996)

¹⁵ On realism in political philosophy see Williams (2005), Geuss (2008) and Galston (2010); compare Stemplowska and Swift (2012). The strand of realism we do not reference here can be understood as an explicit rejection of what it calls ‘political moralism’ (Williams, 2005). The basic idea here is that political philosophy should not consist in the application of moral principles to political situations because political situations generate normativity of their own.

¹⁶ Rousseau’s answer was similar to that of Rawls: taking people as they are is not such a tough constraint because people are malleable and will comply with just requirements of just institutions. The tricky bit is to get to the situation where those just institutions can begin to shape the malleable people.

and Charles Mills (2005), a theory that assumed that people could do the things they cannot do would be an idealising theory, and good theory should not idealise. Partial compliance theory, on this view, could still abstract away from some difficulties (perhaps it could focus on the average capacities and put aside their range), but it should not idealise.

But suppose that the average person was unable to eliminate all their racial biases.¹⁷ Couldn't moral theory still require them not to be racist? On one view, for moral requirements to apply to a person they must have the ultimate control over whether to fulfil them. A person must always be able to avoid condemnation by choosing well. On another view, moral requirements communicate when a person's failure to fulfil them licenses moral criticism. Some think that being racist merits condemnation even if one's attitudes are beyond one's control.

5. The limits of partial compliance theory

Given that partial compliance is inevitable, should not pretty much all moral theory be partial compliance theory? This question is related to what has come to be known as the ideal vs non-ideal theory debate in political philosophy. Rawls divided his own theory of justice into ideal and non-ideal theory. 'Ideal theory' formulates its requirements of justice under conditions of full or near full compliance.¹⁸ By contrast, non-ideal theory deals with the problem of non-compliance (Rawls, 2001: 13). What is the relationship between the two types of theory? Is full compliance theory needed at all? If it is needed, is it needed to engage in partial compliance theory? In what follows we will focus on two possible roles of full compliance theory: epistemic insight and action-guidance.

The epistemic role of full compliance theory is about determining what the requirements of morality and justice are. What morality or justice requires is sometimes obvious, but frequently it is not obvious at all. For example, should there be women only shortlists for selection for a Member of the Parliament? Should economic sanctions be applied to states that attack other states? Should academics accept being paid more than nurses?

One epistemic reason to assume full compliance in our moral theorizing is that moral requirements concern what people should do, not what they will do. And knowing what people would be required to do under conditions of full compliance can help to eliminate self-serving biases.

The exercise quickly gets complicated as readers of Immanuel Kant and Rawls well know. But theorising about what justice requires under conditions of full compliance could nevertheless be epistemically useful for partial compliance theory. Even if the requirements, when issued, do not presuppose the fact of full compliance by all, it is helpful to know what justice would have required if everyone did their part. There are many different things we could do with the information we get when assuming full compliance. When it is understood as an epistemic resource, we can see the information

¹⁷ We are not saying that this is true of the average person.

¹⁸ Rawls mentions also 'favourable conditions' - historical, cultural, social and economic conditions that bear on whether the ideal option is possible (Rawls 2001: 47). For a discussion see Stemplowska and Swift (2013).

as simply a data point to be taken into account when formulating requirements in partial compliance theory (Estlund 2019).

Rawls has suggested that partial compliance theory needs full compliance theory for two specific reasons that go beyond using full compliance as an epistemic resource.¹⁹ These generally concern action-guidance. First, partial compliance theory needs full compliance theory, according to Rawls, to help it identify what is more or less morally urgent here and now under conditions of partial compliance. This is the “urgency role” of full compliance theory (Rawls, 1999a: 267). Second, full compliance theory is needed, on Rawls’s (1999b: 90) account, to orient the requirements of partial compliance theory towards a target. Both suggested roles are controversial. Many have rejected those ideas arguing that partial compliance theory does not need full compliance theory.²⁰

Consider the urgency role. Rawls suggested that the most urgent injustices are ‘identified by the extent of the deviation from perfect justice’ (1999a: 216) and that such injustices should be ‘dealt with first’ (1999a: 267). On this view, partial compliance theory must recommend what ought to be done in light of the requirements of full compliance theory.

But as many have pointed out it is difficult to interpret this claim about priority. Consider what is known in economics as the theorem of the second best. It argues that for each new set of circumstances we need to ask independently what would be best. Following David Estlund (2019: 271–87), here is a simple example to illustrate the idea. Suppose your medical condition (pain) requires you to take pill A to cure it. Pill A, however, has side-effects (severe nausea) that are even more unpleasant than your original medical condition. These side-effects can be neutralized if you take pill B. Thus, to solve your medical condition to your satisfaction you must take both pill A and pill B. Suppose next, however, that pill B becomes unavailable. The theorem of the second best tells us that in such circumstances – when the optimal solution (taking both pills) is not available due to some constraints (the unavailability of pill B) – we should not assume that the second best solution will involve accepting the constraint but otherwise acting the same way we would if the optimal solution were still available, e.g., accepting that we cannot take pill B but proceeding to take pill A. Rather, it tells us that when we cannot take pill B, we should not assume that taking pill A is optimal in our new circumstances.

In fact, Rawls himself cautioned that his principles of justice – and in particular the absolute priority of equal liberty over the difference principle – apply only once societies reach a certain level of economic development (Rawls 1999a). He also wrote of the urgency role of ideal (full compliance) theory that ‘[o]f course, this idea is extremely rough. The measure of departures from the ideal is left importantly to intuition’ (1999a: 216), adding that ‘[i]n the more extreme and tangled instances of nonideal theory...we may be able to find no satisfactory answer at all’ regarding what guidance ideal theory could offer us here (1999a, 267).

The fact that the debate over ideal and nonideal theory arose – at least in its current form – out of the deontological work of Rawls is not a mere accident. Simple consequentialists, such as utilitarians, may insist that we must always get as close as

¹⁹ Some of the discussion draws on Stemplowska (2017).

²⁰ Mills (2005), Farrelly (2007), Sen (2009), Wiens (2015), Gaus (2016), Barrett (2023). For defences of ideal theory see Stemplowska (2008), Valentini (2008), Simmons (2010), Estlund (2019), Táíwò (forthcoming).

possible to total or average utility: the problem of the second best does not arise with respect to whether to do that. But even consequentialists would then need to tell us whether a given mechanism for increasing utility would function with full and with partial compliance. Thinking that the requirements of full compliance theory could simply be adopted by partial compliance theory would land us in trouble.²¹

We will get to Rawls's target role for full compliance theory momentarily. Before doing so, let us consider some other ways full compliance theory has been said to be necessary, or at least important, for partial compliance theory. Theorists such as Christine Korsgaard (1996), Tamar Shapiro (2003), Robert S. Taylor (2009), and Samuel Freeman (2009) have suggested that ideal theory may supplement the guidance partial compliance theory provides (but see also Jubb, 2012). Their suggestion is that full compliance theory offers us "regulative values" that we should follow, even when we find ourselves in the more tangled cases of nonideal circumstances. Korsgaard, for example, has suggested that when citizens face injustice, "civil disobedience is better than resorting to violence not just because violence is bad in itself, but because of the way in which civil disobedience expresses the democratic principles of the just society it aspires to bring about" (Korsgaard, 1996: 148, quoted in Taylor, 2009: 489-90).

It's not clear, however, that we need full compliance theory to identify these regulative values for nonideal theory. Both partial and full compliance theory can draw on a theory of ideals (Hamlin and Stemplowska 2012). A theory of ideals is concerned with identifying and specifying our ideals (and values), whereas partial and full compliance theory focus on which actions, practices, or institutions might best realize them. This is, of course, a porous distinction: specifying a value may require reflecting on what it would take to realize it. Nonetheless, the distinction is useful because it allows us to see that theorizing about ideals is not the exclusive province of full compliance theory. This explains why partial compliance theory can try to claim independence from full compliance theory: it does not need the latter to be able to include values and ideals in its analysis of what ought to be done.

A different way of thinking about the relevance of full compliance theory for partial compliance theory is that it could offer relevant benchmarks for some questions. For instance, partial compliance theory may need to consider what requirements would be recommended by full compliance theory to identify the costs of noncompliance. In order to know how "bad" or "costly" noncompliance is, we need a benchmark to compare it with. Full compliance theory can play this role.

For example, suppose that due to noncompliance only some very slow climate change reforms can be undertaken, leaving some adversely affected. The actual cost of this can be grasped by comparing a situation where everyone complied with the requirements at hand – this is the benchmark – to what the situation looks like when some people fail to comply. Of course, knowing the costs of noncompliance does not by itself tell us who should bear them here and now or how urgent it is to mitigate for them. Many other factors are at play. Nonetheless, the costs of noncompliance would seem to give those who are poorly situated due to it an extra reason to object to their situation.

What to make of this reason? As argued in rabid dog case, the complaint that one's situation is due to noncompliance may make no difference to whether one gets priority over those whose situation is not due to noncompliance. But, as we also mentioned

²¹ See also Swift, 2008: 375-378.

there, if the noncompliance is of a certain type, it may make a difference to who should bear the costs of redress. In some cases, victims of noncompliance may have a right to pursue compensation claims against noncompliers that they would otherwise lack.

The suggestion that full compliance theory can offer a benchmark for some partial compliance questions goes against a thesis advanced by Amartya Sen (2009). Sen has argued that normative theorizing is in need of what he calls “comparative theory” as opposed to “transcendental theory.” Transcendental theory concerns the ideal end point of society (listing its institutional arrangements). Comparative theory asks what might constitute a feasible improvement from where we are here and now. Putting aside the possibility that where we are here and now is at or near the ideal point, comparative theory is a version of partial compliance theory and transcendental theory is a version of full compliance theory. Sen’s main point is that comparative (partial compliance) theory does not need transcendental (full compliance) theory. To illustrate his point, Sen observes that in order to compare a Dalí and a Picasso, we don’t need to know that Da Vinci’s the *Mona Lisa* is the best painting: we can engage in comparisons without knowledge of the ideal (2009: 16). Sen’s overall conclusion is that full compliance theory is not only useless to partial compliance theory, it is an unhelpful distraction.

Sen might, then, respond to our suggestion that full compliance theory can offer a useful benchmark, for instance, for calculating costs, by appealing to the problem of the 2nd best. Even if it is true that in ideal circumstances with full compliance Anna would be required to do X, it does not follow that Anna should be made to bear the costs of her failure. This is because, as the theorem of the second best shows, we cannot simply assume that getting closer (in one respect) to the full compliance outcome by pursuing compensation from Anna is best in the nonideal circumstances in which the ideal outcome is not accessible.

Our suggestion, however, is not that compensation must be sought but that partial compliance theory should consider whether it should be sought and from whom, and knowledge of who has not complied would be relevant here. That is, even if we ultimately decide that compensation is not appropriate, we still need to ask whether it is. We may think that while compensation would not have been appropriate or possible in the past, it could be now. In short, we should recognize that complete (and correct) partial compliance theory should seek to theorize in light of the findings of full compliance theory in order to make its requirements more accurate.

What of Rawls’s second suggestion that partial compliance theory needs full compliance theory to set the target for its suggested social reforms? (Rawls 1999b: 90). Suppose that we are currently facing two possible (but politically contested) short-term reforms to improve our state educational system. One reform encourages private schools to partner with state schools as a way of improving the latter’s access to resources. The other encourages state schools to spend less on infrastructure as a way of freeing resources for better teachers. Suppose both reforms would plausibly increase children’s access to good education. Since this is an important value for any educational system, this would constitute an improvement.

Even though both reforms would be improvements over the status quo, we still need some way of deciding between them. One way of doing so would be by recognising that we have many different views about several aspects of the ideal educational system. And if we know where we ideally want to end up, this can help us select a reform. If we think that in an ideal system, there would be no private schools, the first reform is less likely

to lead us there. If we think that in the ideal system, private education would be allowed but teachers would self-select into the profession without (much) regard for salary, we may think that the second reform is less likely to lead us there. The ultimate decision is of course more complicated since neither reform advances the achievement of either conception of our ideal of education. But at the very least it would seem odd not to incorporate their likely and possible effects on our ability to reach the target.

Some have questioned whether theorising about full compliance is really needed to set the target. We take it that the debate is best understood as a debate about the reach and complexity that full compliance theory needs to exhibit to qualify as full compliance theory. There are different answers available. Suppose, for instance, that any theory that at some point considers the implications of a given requirement under conditions of full compliance would qualify as full compliance theory. We might call this *narrow* full compliance theory. Full compliance theory in this narrow sense is an essential aspect of partial compliance theory. A partial compliance theory needs to be able to consider what might happen if people act better than what the theory expects of them. This includes narrowly considering what would happen if everyone did follow a given requirement. Because, after all, they might just do it.²²

But full compliance theory could also be understood in a *comprehensive* sense. Interpreted this way, full compliance theory is a theory that asks what *all* the key aspects of society would look like under full compliance.²³ Rather than picturing what would happen if each person followed a particular requirement, full compliance theory understood in the broad sense entails looking at what would happen if each person complies with each available requirement.

It is less clear why comprehensive full compliance theory is necessary for setting the target for a partial compliance theory, because we may be unable to ever reach it.

At the same time, comprehensive full compliance theory might simply be the wrong view. Even Rawls's theory hardly offers a fleshed out vision of a truly just society. Other contemporary theorists do even less of that, so full compliance theory in the comprehensive sense may not be on offer in any case. By contrast, if full compliance theory is understood in the narrow sense as a theory that at some point considers what a requirement would deliver if all complied, then this type of theory seems inescapable for those who wish to offer good partial compliance theory.

6. Conclusion

Systematic reflection on what follows from the fact that people often fail to do what is required of them can help us see more clearly what morality requires. The fact that people only partially complies with requirements of morality also raises numerous distinct issues for moral and political philosophy.

In this chapter, we have explored partial compliance theory. We have suggested that there is a close link between injustice and noncompliance, assessed how theorists have grappled with the question of what compliers can be expected to do when others fail to comply, which requirements can be issued given that people will often fail to comply,

²² It also needs to reflect on the optimal destination for a given reform.

²³ There can of course also be intermediate versions available, where full compliance theory considers compliance with some of the requirements that apply.

and whether it is possible to do partial compliance theory without theorising about full compliance as well.

7. References

Abdel-Nour, Farid (2018). 'Responsibility for structural injustice', *Ethics and Global Politics* 11(1), 13-21.

Arneson, Richard (2004). 'Moral Limits on the Demands of Beneficence?' In Deen K. Chatterjee (ed.). *The Ethics of Assistance*. Cambridge: Cambridge University Press, 33-58.

Bajaj, Sameer (2019). 'The weight of fairness', *Politics, Philosophy and Economics* 18(4), 386-402.

Barrett, Jacob. (2023). 'Deviating from the ideal.' *Philosophy and Phenomenological Research*, 107(1), 31-52.

Barry, Christian and Kate Macdonald (2016). 'How should we conceive of individual consumer responsibility to address labour injustices?', in Yossi Dahan, Hanna Lerner, and Faina Milman-Sivan (eds.), *Global Justice and International Labour Rights*. Cambridge: Cambridge University Press, 96-97

Butt, Daniel. (2019). 'Restitution post bellum: property, inheritance, and corrective justice', *Journal of Applied Philosophy* 36(3), 357-365.

Butt, Daniel. (2021). 'What structural injustice theory leaves out', *Ethical Theory and Moral Practice*, 24, 1161-75.

Cohen, G.A. (1997). 'Where the action is: On the Site of distributive justice', *Philosophy & Public Affairs* 26: 3-30.

Cohen, G. A. (2008). *Rescuing Justice and Equality*. Harvard University Press.

Cohen, Jonathan L. (1981). Who is starving whom? *Theoria* 47 (2): 65-81.

Cullity, Garrett. (2006). *The Moral Demands of Affluence*. Oxford: Oxford University Press.

Kamm, Frances .M. (2000). 'Does Distance Matter Morally to the Duty to Rescue?', *Law and Philosophy* 19, 655-681.

Dietz, Alexander (2016) 'What we together ought to do', *Ethics*, 126, 955-982.

Estlund, David. (2019), *Utopophobia*. Princeton University Press.

- Farrelly, Colin. (2007). 'Justice in Ideal Theory: A Refutation.' *Political Studies* 55(4), 844-64.
- Freeman, Samuel (2009). *Justice and the social contract: Essays on Rawlsian political Philosophy*. Oxford University Press.
- Galston, William. (2010), 'Realism in Political Theory', *European Journal of Political Theory* 9, 385-411.
- Gaus, Gerald. (2016). *The Tyranny of the Ideal: Justice in a Diverse Society*. Princeton: Princeton University Press.
- Geuss, Raymond. (2008), *Philosophy and Real Politics*. Princeton University Press.
- Goodin, Robert and Christian Barry (2021). 'Responsibility for structural injustice: A third thought', *Politics, Philosophy, and Economics*, 20(4), 339-356.
- Hamlin, Alan and Zofia Stemplowska. (2012), 'Ideal Theory, Nonideal Theory and the Theory of Ideals', *Political Studies Review*, 10: 48-62.
- Hoesch, Matthias. (2018). 'Taking up the Slack' in the Context of Refugee Protection: Remarks on David Owen. *ZEMO, Zeitschrift für Ethik und Moralphilosophie* 1: 163-175.
- Jubb, Robert. (2012). 'Tragedies of Non-ideal Theory', *European Journal of Political Theory*, 11, 229-246.
- Karnein, Anja. (2014). Putting Fairness in its Place: Why There is a Duty to Take Up the Slack. *Journal of Philosophy* 111: 593-607.
- Korsgaard, Christine. (1996), *Creating the Kingdom of Ends*. Cambridge University Press.
- McKeown, Maeve. (2021). 'Structural Injustice', *Philosophy Compass*, 16(7), 1-14.
- Miller, David. (2011). Taking up the slack? Responsibility and justice in situations of partial compliance. In Carl Knight and Zofia Stemplowska (eds.). *Responsibility and Distributive Justice*. Oxford University Press.
- Miller, David. (2016). *Strangers in Our Midst*. Harvard University Press.
- Mills, Charles W. (2005). "'Ideal Theory" as Ideology', *Hypatia* 20, 165-84.
- Murphy, Liam. (2000). *Moral Demands in Nonideal Theory*. Oxford University Press.
- Nuti, Alasia. (2019). *Injustices and the Reproduction of History*. Cambridge University Press 2019.

- O'Neill, Onora. (1996). *Towards Justice and Virtue*, Cambridge University Press.
- O'Neill, Onora. (1988) 'Abstraction, Idealization and Ideology in Ethics' in J.D.G. Evans (ed.), *Moral Philosophy and Contemporary Problems*, Cambridge University Press, 55-69.
- Overland, Gerhard. and Christian Barry (2011). 'Do Democratic Societies Have a Right to Do Wrong?,' *Journal of Social Philosophy* 42: 111–31.
- Owen, David. (2016). 'Refugees, Fairness and Taking up the Slack: On Justice and the International Refugee Regime', *Moral Philosophy and Politics*, 3(2), 141–164.
- Owen, David. (2018). 'Taking up the Slack' in the Context of Refugee Protection: A Reply to Matthias Hoesch. *ZEMO, Zeitschrift für Ethik und Moralphilosophie* 1, 177–184.
- Parekh, Serena. (2017). 'Feminism, Structural Injustice, and Responsibility' in Ann Garry, Serene J. Khader, Alison Stone (eds.). *The Routledge Companion to Feminist Philosophy*. Routledge.
- Parfit, Derek. (2011). *On What Matters: Volume 1*. Oxford: Oxford University Press.
- Rawls, John. (1999a). *A Theory of Justice*, revised ed. Oxford University Press.
- Rawls, John. (1999b). *The Law of Peoples*. Harvard University Press.
- Rawls, John. (2001). *Justice as Fairness: A Restatement*, Erin Kelly (ed.). Harvard University Press.
- Rousseau, Jean-Jacques. (1762/2002). 'The social contract.' In *The Social Contract and the First and Second Discourses*. Susan Dunn (ed.). Yale University Press.
- Reiman, Jeffrey. (2012). 'The Structure of Structural Injustice: Thoughts on Iris Marion Young's "Responsibility for Justice"', *Social Theory and Practice*, 38(4), 738-751
- Sangiovanni, Andrea. (2018). 'Structural Injustice and Individual Responsibility', *Journal of Social Philosophy*, 49, 461-483
- Schwenkenbecher, Anne. (2021). *Getting Our Act Together*, Routledge.
- Sen, Amartya. (2009). *The Idea of Justice*. Allen Lane.
- Shapiro, Tamar. (2003). "Compliance, Complicity, and the Nature of Nonideal Conditions", *Journal of Philosophy*, 100: 329-355.
- Simmons, A.J. (2010). 'Ideal and Nonideal Theory', *Philosophy and Public Affairs*, 38, 5-36.

- Singer, Peter. (1972). Famine, Affluence, and Morality. *Philosophy & Public Affairs* 1(3), 229–243.
- Singer, Peter. (2002). *One World*. New Haven, CT: Yale University Press.
- Stemplowska, Zofia. (2008), 'What's Ideal About Ideal Theory?', *Social Theory and Practice*, 34, 319-340.
- Stemplowska, Zofia. (2016). 'Doing More Than One's Fair Share.' *Critical Review of International Social and Political Philosophy*, 19(5), 591-608.
- Stemplowska, Zofia. (2017). 'Nonideal Theory' in Kasper Lippert-Rasmussen, Kimberly Brownlee and David Coady (eds.), *A Companion to Applied Philosophy*, Wiley-Blackwell.
- Stemplowska, Zofia. (2019). 'Coercing compliers to do more than one's fair share.' *Zeitschrift für Ethik und Moralphilosophie*, 2(1), 147-160.
- Stemplowska, Zofia and Adam Swift. (2012). 'Ideal and Nonideal Theory,' in David Estlund (ed), *The Oxford Handbook of Political Philosophy*. Oxford University Press.
- Stemplowska, Zofia and Adam Swift. (2013), 'Rawls on Ideal and Nonideal Theory,' in Jon Mandle and David Reidy (eds.). *A Companion to Rawls*. Wiley-Blackwell.
- Swift, Adam. (2008), 'The Value of Philosophy in Nonideal Circumstances', *Social Theory and Practice*, 34, 363-387.
- Táíwò, Olúfémi O. (Forthcoming). 'Uncommon Features: Defending Ideal Theory with Model-to-World Inference.' *Routledge Handbook on African Political Philosophy*. Routledge.
- Táíwò, Olúfémi O. (2022). *Reconsidering Reparations*, Oxford University Press.
- Taylor, Robert S. (2009), 'Rawlsian Affirmative Action', *Ethics*, 119: 476-506.
- Unger, Peter. (1996). *Living high and letting die*. Oxford: Oxford University Press.
- Valentini, Laura. (2008). 'On the apparent paradox of ideal theory.' *Journal of Political Philosophy*, 17(3), 332–355.
- Wiens, David. (2015). 'Against Ideal Guidance.' *The Journal of Politics*, 77, 433–445.
- Williams, Bernard. (2005). *In The Beginning Was The Deed: Realism and Moralism in Political Argument*. Princeton University Press.
- Young, Iris Marion. (2011). *Responsibility for Justice*. Oxford University Press.