States and the Limits of Feasibility:

Enforcement, Morality and Possibility

Daniel Guillery

University College London

Thesis to be submitted for the degree of PhD

I, Daniel Guillery, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Signed:

Abstract

Enforcement seems to be an essential and ubiquitous feature of state societies. My thesis explores arguments for the kind of general and exclusive moral permission to enforce that states claim, and in particular the role that feasibility considerations play in them. I argue that premises about the infeasibility of alternatives to a state's enforcement are essential to the success of any such argument. States' permission to enforce can be justified, if at all, in response to the unfortunate circumstances in which we find ourselves. I develop a general *multivocal* account of the concept of feasibility, according to which the concept can be made precise in many different ways, no single one of which is obviously privileged as uniquely relevant to moral theory. This account has the result of casting doubt on the assumption that states' permission to enforce can be taken for granted. Arguments for this permission may succeed when we make their feasibility premises precise in some ways, but not others. Understanding this, I argue, helps illuminate how we ought to think about and treat the state enforcement we face in the real world.

Impact Statement

This thesis investigates the moral permissibility of state enforcement. Its arguments, if accepted, will have a quite significant impact on both the discipline of political philosophy and on the practice of politics. My arguments show that states' claimed general and exclusive moral permission to enforce compliance with their commands or directives should not be taken for granted, as it tends to be. Firstly, this has an impact on how political philosophers ought to think about state legitimacy. While it is fairly common among political philosophers to be sceptical about general *obligations* to obey the state, it is usually assumed that states' general and exclusive moral permission to enforce (which I refer to as 'legitimacy') is on safe ground. My conclusion, then, ought to encourage political philosophers to explore further the kinds of arguments that are available for state legitimacy in this sense, and to ask whether or not we should in fact consider our states to be acting morally permissibly. More specifically, I argue that plausible arguments for state legitimacy depend on facts about feasibility, and that it is not obvious that the necessary feasibility facts hold. This establishes a need to explore exactly what feasibility claims need to be true in order for existing states to be legitimate, and whether these feasibility claims in fact are true. More generally, my arguments demonstrate the importance of unrealistic political theory to thinking about state legitimacy, where previously debates concerning realisticness of theory had focused on justice.

It also follows that it is only insofar as it is not feasible to do better that our states can be justified in acting as exclusive enforcers. This has the consequence that we should think about our states as legitimate, if at all, only in response to unfortunate circumstances, and that a world without these states is in some sense *ideal*. Although this need not mean that we ought to see to abolish our states, it does have an impact on how we ought to treat our states' enforcement here and now. We should seek to minimise it where possible, consistently with other weighty values or moral requirements, and if and when we ought to accept it, this is a tragic fact, not something to be welcomed.

The thesis also contains a rigorous account of the concept of feasibility, which feeds into the argument concerning state legitimacy, but also stands independently of it and, if accepted, should have a significant impact on how political philosophers think about the relation between feasibility and moral and political theory. It tends to be thought that proposals in moral or normative

political theory are straightforwardly ruled out if *infeasible*. My account of this concept, though, has the consequence that there is no single straightforward feasibility constraint of this sort on moral and political theory (or at least, it should not be assumed that there is). Thus, it will not suffice for rejecting a theory to declare that what it calls for is infeasible.

Contents

Acknowledgments

Thanks are due to many who have shaped the content of this thesis and supported me through its creation. Firstly, I am indebted to the UCL philosophy department as a whole for having provided an ideal and stimulating environment in which to do this work, and to the London Arts and Humanities Partnership for having funded it. Its shape and argument owes a good deal to the questions and comments of many who have heard me present earlier versions of various chapters, including participants in the 2014 Brave New World Conference at the University of Manchester; the 6th Meetings on Ethics and Political Philosophy at the University of the Minho; the 2016 Pavia Graduate Conference in Political Philosophy at the Università degli Studi di Pavia; the 2017 Warwick Graduate Conference in Legal and Political Theory; the 2017 Freedom and Autonomy Conference at Birkbeck College; the UCL MPhil Stud thesis preparation seminar, PhD Work in Progress seminar and moral and political philosophy work in progress group; and the Brown Graduate Political Philosophy Workshop. Thanks also for discussion of parts of this work, or ideas that in some form or other found their way into it, to Hannah Carnegy-Arbuthnott, Ben Fardell, Rob Joynt, Fiona Leigh, Michael Martin, Rowan Mellor and José Zalabardo. Special thanks are owed to those who have read and discussed with me in great detail large sections of the work that later became this thesis: Han van Wietmarschen, Dave Estlund and especially my supervisor, Amanda Greene. Special thanks go also to my father and to Ella, for proof-reading and constant support. Finally, above all, thanks are due to my primary supervisor, Jo Wolff, whose contributions to this thesis are too extensive to enumerate. It is thanks to his tireless work and support (even despite a change of institution) that the following pages contain a thesis at all.

Introduction

It is sometimes thought that questions about how states and their subjects morally relate to each other are the *fundamental* questions of political philosophy. If 'political society' is understood as society organised by a state, then the motivation for this is obvious: these questions concern the moral status of the *subject matter* of political philosophy. I do not think that this makes these questions fundamental in any deep sense. If political philosophy is defined thus, there is no reason to suppose that it in turn is particularly fundamental to normative *social* philosophy broadly (the study of the normative principles, ideals and values relevant to large-scale questions about how *societies* ought to be structured and organised). Societies need not be structured by states, and so there are other questions in normative social philosophy that are not necessarily subsequent to those about states.

Nevertheless, it cannot be denied that the moral evaluation of states is of great importance for the morally concerned thinker in the world in which we live.[1] We live on a planet whose habitable surface is entirely within the claimed jurisdiction of some state or other. Almost the entirety of this is also under the de facto control, the enforcement power, of some state or other. Though there *can* (and perhaps should) be non-state forms of society, they are barely existent in the modern world. Given the ubiquity of states, then, and the extent to which they shape our lives, it should not be surprising that much philosophical effort has been devoted to questions about state and subject. It is important to think about what moral attitude we should take to these states that surround us and the demands that they make of us. This thesis aims to contribute to this task by addressing one particular dimension of the moral evaluation of states. This is, roughly, states'

---

[1] There is a growing 'realist' trend in political theory, some variants of which reject the application of moral principles or values to the political realm (for surveys of the realist trend see Galston (2010) and Rossi and Sleat (2014) and for criticism see Estlund (2017c)). I do not find this view attractive, since the political realm, just like all other social realms, is constituted by human activity, about which we can ask moral questions. We ought to be able to ask questions, at least, about whether political society is something we should *have* in the first place, or rather should seek to avoid; these, presumably, cannot be settled by values or principles internal to the political. In any case, this thesis will not address such views. I assume that political society is an appropriate target for moral evaluation, and what follows will be an exploration of the *moral* status of an aspect of states. (Some others claim that the kind of moral principles that apply to politics are distinctive and sui generis, (see, e.g. Rossi (2012)). It seems likely that there will be certain sorts of moral principle that are particular to the political realm, but there seems no *ex ante* reason to expect that general moral principles or considerations will *not* hold in the political domain, so I will not begin by assuming that they do not.)

possession of a general moral permission to enforce compliance with their demands, which I will stipulatively refer to as state legitimacy. (I will attempt to make this notion more precise shortly.) This thesis hopes to cast some doubt on the assumption often made that state legitimacy, in this sense, is on safe ground (for at least some actually existing states). It will do this by way of an examination of the importance of *feasibility* considerations to arguments for state legitimacy.

1. Anarchism

One way to understand this thesis is as a call for the *anarchist* challenge (or a form of anarchist challenge) to be taken seriously. Anarchism is primarily some form of *rejection* of the state, and since the world as we currently know it is entirely occupied by states, it is a rejection of the status quo. There are, though, many forms of anarchism, and a number of different theses which might be taken as definitive of it. Many of these have in common a call for something radically different from the status quo, or some form of rejection of something the eradication of which would be a radical departure from the status quo.

It is very common for anarchism of various forms to be rejected on grounds of *infeasibility*.[2] What it calls for is *infeasible*, it is argued, and so a political theory recommending it must be wrong. Or, there is no feasible alternative to what it claims is impermissible, thus it must be mistaken about this impermissibility. Or, it is wrong that a certain aspect of the status quo is morally bad or corrupt, because there is no feasible alternative. (In fact, the feasibility critiques made of anarchism are not usually so straightforward: it is not usually claimed simply that what anarchism calls for is infeasible, nor that there is no feasible alternative to what anarchism rejects. Rather, it is usually claimed that there are certain particularly weighty values or principles whose realisation in conjunction with what anarchism calls for is infeasible. In other words, what anarchism calls for is not feasible *in a desirable way*.)

Given the radical nature of anarchism, the prevalence of such feasibility critiques is not unexpected. However, there is surprisingly little understanding of the concept of feasibility, and of the role it plays in constraining moral and political theory.[3] Thus, it is not clear when the truth of certain feasibility facts warrants

---

[2] These sorts of feasibility critiques can be found in, for example, Miller (1984) and Wellman (1996) 232.

[3] There has been a recent increase in interest in the concept, and a few sophisticated accounts have emerged. However, I will argue in chapter 5 that none of these is fully successful.

the rejection of some moral or political theory. The focus of my thesis, then, will be on the role of feasibility considerations in arguments attempting to overcome the anarchist challenge. I will argue that arguments which might successfully establish state legitimacy (for some existing states) depend on feasibility premises. In the final part of the thesis, I will offer an account of the concept of feasibility which has the consequence that arguments relying on feasibility premises should not be *assumed* to work. Further consideration is needed of these feasibility premises and whether they support the arguments in which they are used.

What precisely, though, is anarchism? The term is a contested and multifarious one. For one thing, the term is claimed by certain political movements, and so is sometimes used to refer to an ideology, a *collection* of claims. It is also sometimes used simply to designate political troublemakers or similar. Setting aside these uses, the term centrally refers to some form of rejection of the state.[4] Since there are a number of different dimensions of evaluation of states, there are a number of different normative theses that might be labelled 'anarchist'.

Let me distinguish three key sorts of question about the moral status of states, negative answers to all of which might be labelled 'the anarchist thesis'. All three questions are independent of each other: for any given state, some of the questions might receive positive answers while others receive negative answers. The first sort of question asks whether the *existence* of the state in question is, in some sense *a good thing* or *morally acceptable.* This might be termed the question of whether the state is *justified.* (This is, I think, what A. John Simmons has in mind when he famously distinguishes justification from legitimacy, a term which he uses differently from me.)[5]

The word 'justification' is often used in connection with actions, as when we ask whether someone is *justified in doing* something. This is not the way I mean to use the term here. Existing is not an action, and entities cannot be *justified in existing* in this sense. It may be that when people talk about *justifying the state*, or showing that a state is *justified*, they mean to refer elliptically to showing that a state is *justified in doing* certain things (perhaps those things essential to or characteristic of states). Since enforcement is often taken (rightly, as I will suggest below) to be

---

[4] When the term is used to refer to a single normative thesis, it does not *always* describe a rejection of the *state*: sometimes it is used to refer to a rejection of something else, such as authority, coercion, power, domination or hierarchy in general, or to a positive call for something, such as decentralisation, but let us set aside these uses as well. For these, see for example Clark (1978), Kropotkin (1970) 284, Malatesta (quoted in Clark (1978) 5 and Berman (1972) 28) and Woodcock (1963).

[5] Simmons (2001) 122-57

essential to states, showing that a state is *justified in doing* things essential to its statehood may end up amounting to establishing what I above called its *legitimacy* (which I will turn to below). But I think that talk of states being justified may also refer to something different, having to do with the moral status of the state's *existence*. When I refer to a state's *justification* or being *justified*, I refer exclusively to this latter thing.

As we will see in chapter 2, there are really a number of slightly different questions here (a number of different properties of states that we might be asking about). But these different questions seem to form a family, having to do with the quality of the state and whether it is something whose existence we should accept or be glad for. Justification anarchism, then, would be something like the view that it would be better if we did not have a state, or that we ought to get rid of our states. (Simmons points out that there is a distinction between a priori and a posteriori anarchism, and such a distinction can also be made within justification anarchism: a priori justification anarchism holds that all states are necessarily unjustified, while a posteriori justification anarchism holds only that all currently existing states are unjustified.[6] An equivalent distinction can be made also for the other anarchist views about to be discussed.)

A second question, which has been granted a great deal of attention by political philosophers in recent years, is the question whether the state has obligating power over its subjects in the following sense:

> *Obligating Power*. A state X has obligating power over an individual Y in some domain D if and only if X's commands/directives/laws in D create (at least pro tanto) duties/obligations/pre-emptive reasons for Y to act as commanded/directed, independent of their content (within some limits).[7]

This is, I think, the same question as the traditional question of political obligation. To say that subjects of a state (or certain subjects of a state) have a general obligation to obey whatever laws or directives that state makes just

---

[6] Simmons (1996) 105

[7] On pre-emptive reasons see Raz (1988) ch. 3 and (1979) 18; a pre-emptive reason to φ is a reason that pre-empts, that is, excludes from consideration or replaces certain other reasons for or against φ-ing. X's having obligating power over Y involves (in Hohfeldian terms) it having a *moral power* to change Y's normative situation, but it is consistent with either X's also having a *claim right* to Y's obedience or not. For discussion of this distinction see Copp (1999) 10-21. For X's commands/directives to *create* a duty (etc.) need not require that X's command be in any sense the *ultimate* source of the duty. For instance, if Y promises to obey whatever commands X makes, and then on some subsequent occasion X commands Y to φ, Y now has a duty to φ that she did not have before. X has *created* this duty, even though, in some sense, its ultimate source is Y's promise, not X's command; Y only has the duty to φ as a part of the more general duty to do whatever X commands.

because they are laws or directives the state has made, independent of their content, is just to say that that state has the power (or capacity) to make its subjects (or certain of its subjects) obligated to obey its laws or directives just by making one.[8] (I introduced this question as the question of obligating power, rather than as the question of political obligation because I am interested in the moral evaluation of *states*, rather than that of citizens.) Here, the anarchist view (obligating power anarchism) is the denial of states' obligating power, again either a priori or a posteriori: either the claim that no state could possibly have obligating power over any (or all) of its subjects, or the claim that no existing state in fact has obligating power over all of its subjects.

The third question I will call, stipulatively, the question of *state legitimacy*. This is the question whether the state in question has a general permission to enforce compliance with its commands or directives. (I use 'enforce' in a way to be explained in more detail below, which may not correspond perfectly with the ordinary use: roughly, to enforce is to make another do something whether or not they want to.) The anarchist view here is that states lack any general permission to enforce compliance with their commands or directives. Again, this could be a priori or a posteriori: either the claim that no state could possibly have such a general permission, or simply that no existing state has such a permission. It is this question in which I will be interested in this thesis, and it is the *a posteriori* anarchist claim that I argue should be taken seriously. I will go on to refine the question below. For now, though, let me point out that it is independent, at least conceptually, of the other two questions.[9] First, it does not *immediately* follow from a state's being justified in the above sort of sense (roughly, its existence being a good thing or morally acceptable) that it is permitted to enforce its directives over its subjects. The acceptability or goodness of a state's *existence* and what it is permitted to do are two separate things. However, it might well be thought that legitimacy *does* follow from justification, even if it is not a straightforward conceptual entailment, and I think something like this thought may underlie the extent to which legitimacy is sometimes taken for granted. This line of thought will be the subject of chapter 2.

---

[8] X's having content-independent obligating power over Y can be understood as it being the case that, for any φ, *that X commanded it* is sufficient for, and the reason for, Y's being required to φ *so long as* some conditions are met (and these conditions give the limits mentioned in the definition above).

[9] I also think that justification and obligating power are independent of each other, and I take this to be the main point Simmons (2001) makes when he distinguishes justification from legitimacy. He uses the term 'legitimacy' to mean roughly the conjunction of what I call obligating power and legitimacy, and I think his arguments primarily address the independence of obligating power from justification.

Legitimacy (in my sense) is also independent of obligating power:[10] if a state's citizens do generally have a content-independent obligation to obey its directives, it does not follow from this that it is permissible for the state to *enforce* their doing so; and, on the other hand, if there is no such obligation, it does not follow either that the state is *not* permitted to enforce compliance.[11] In other words, neither obligating power nor legitimacy follows from the other. In general, it does not follow from something's being obligatory for you that it is permissible for anyone to enforce your doing that thing, and it also does not follow from its being permissible for someone to force you to do something that it is obligatory for you to do that thing. For the first, consider promising: it is usually thought that in most cases we have an (at least pro tanto) obligation to keep our promises, but it is not standardly thought that it is usually permissible for anyone to *enforce* people's keeping their promises.[12] Or take a requirement of gratitude. It is quite plausible that if someone does something to benefit us significantly at some cost to themselves (saves us from drowning, advises us on how to achieve some goal, for instance), we are under some obligation to show gratitude in some way. It is not plausible at all in most such cases that it is permissible for anybody to *force* us to show gratitude. For the second, consider a case where I use force to move you out of the way of an oncoming train to which you are oblivious: what I do seems permissible, but there it would surely be wrong to say that it was *obligatory* for you to move out of the way of the train you were unaware of.

While the state case is of course a special one, it seems clear in general that questions of the permissibility of enforcement are *separable* from questions about

---

[10] There is much debate about whether legitimacy is properly analysed such that it entails political obligation (see the next paragraph). My claim here does not take a stand on this question; I use 'legitimacy' here in my stipulative sense and so my claim is the less controversial one that *the general permission to enforce* is independent of obligating power.
[11] Others have noticed that these are separate questions, e.g. Green (1988) 243; Sartorius (1999) 144-6 and Wellman (1996) 212. Patrick Durning (2003) argues persuasively that obligating power and legitimacy are separable, although he focuses on the entailment from legitimacy to obligating power. I think it is just as clear, if not more so, that there is no entailment the other way around. (Both Green (2004) and Wellman seem to think that political obligation *does* entail the legitimacy of state coercion; that is, they think that it is not possible for a state to have obligating power but no permission to enforce. Neither, though, offers any argument for this claim and I can see no reason to think that it must be permissible for someone to enforce a set of rules in order for one to have an obligation to obey those rules.) There are arguments that legitimacy entails obligating power in e.g. Klosko (1992) 38, 45 and Wyckoff (2010) on the basis of a principle that one may only force people to do things that they have a duty to do. These arguments are convincingly refuted by Durning (2003), and my oncoming train example below also provides a counterexample to this principle.
[12] There are views of morality according to which you automatically forfeit certain of your rights by acting morally wrongly (see, for instance, Simmons (1991) sec. 4). I do not find such views plausible, but even if they are, that *some* rights are forfeited by morally wrong behaviour does not suffice to make enforcement automatically permissible.

obligation: answers to one do not *automatically entail* answers to the other, and it is not obvious why the state case should be any different in this respect. It seems plausible that it is *possible* to have an obligation to obey someone's commands if, for example, I promise to obey. (It is unclear whether it is possible for a promise to bind content-independently over a whole lifetime, but it at least seems plausible that if I promise to obey your commands, this generates a pro tanto, content-independent obligation or duty to obey you for some period of time.) There is no reason to suppose, though, that promising to obey in itself grounds any permission to enforce my compliance with your commands.

I have chosen to refer to the question I am interested in as the question of state legitimacy. Before I move on, I want to note that there is a significant segment of the literature on state legitimacy that I will set aside. As Patrick Durning points out, there is one dispute concerning legitimacy that is linguistic or analytical in nature.[13] It concerns what it means to call something a legitimate state: which Hohfeldian rights a state must have in order to be called legitimate.[14] Does a state's having a liberty right to 'act as a state' suffice for its being genuinely *legitimate*? Or is it also necessary to have a claim right to non-interference or a power to create duties in people residing in a certain territory? This, Durning astutely notes, is a debate concerning the proper application of the term 'legitimacy', not a substantive dispute. It is a debate about what the proper analysis of our concept of legitimacy is. I do not have any contribution to make to this debate (and I am not entirely convinced that 'legitimacy' is not too much a philosopher's term of art to admit of a proper conceptual analysis). I mean my restrictive use of the term 'legitimacy' to be purely stipulative. My using it in this way is not wholly idiosyncratic, though: this is a way in which the term is sometimes used, and the general moral permission to enforce that I am calling 'legitimacy' is taken by many authors to be at least a part of what is required for legitimacy.[15] Nevertheless, I do not mean my use of this term as a proposal about the nature of legitimacy. If you think that what I am calling 'legitimacy' is not legitimacy, you can simply substitute some other word.

---

[13] Durning (2003) 374–5

[14] Contributions to this debate can be found in, for example, Applbaum (2010), Copp (1999), Edmundson (1998), Greenawalt (1999), Song (2011), Wyckoff (2010) and Zhu (2017). There is some discussion of the debate in Peter (2012).

[15] The term seems to be used in my way by, e.g. Estlund (2009) and Wellman (1996) and (2001). Green (2004) uses the term in a similar way.

2. *A common assumption*

There is an assumption that I think is quite commonly implicit in political philosophy, if less commonly explicit, which is that legitimacy (in my sense: roughly, states' general permission to enforce compliance with their directives) is on safe ground and can be taken for granted. By that I mean not that it is assumed that states can be supposed always to be legitimate (that is certainly not a common view), but rather that it is assumed that the *existence* of legitimate states (and indeed, the possibility of state legitimacy) can be taken for granted. In other words, it is quite often assumed that the *legitimacy*-anarchist challenge is not a serious danger: it is not the case that all existing states fail to be legitimate in this sense. Similar assumptions are not made with respect to other anarchist challenges: in particular, the *obligation*-anarchist challenge tends to be taken quite seriously. I think it is often supposed, however, that legitimacy is on *safer* ground than obligating power.

There has been a lot of work done attempting to find stable grounds for political obligation. This search has proved difficult and arguments attempting to establish such grounds have tended to be subject to forceful objections.[16] The view that there are no general political obligations (obligations or duties to obey the commands/directives/laws of the state just because they are commands/directives/laws of that state, i.e., independent of their content) has become relatively widespread.[17] It is not so commonly thought that legitimacy of states (*in my sense*) is in similar danger. State legitimacy is much discussed in political philosophy (and although it is not always, the permission to enforce is often taken to be necessary, and sometimes sufficient, for legitimacy). But most often what is asked is exactly what conditions a state must fulfil in order to be legitimate, rather than whether it is *possible* for a state to be legitimate.[18] These accounts do not generally address the legitimacy anarchist (who believes that no states are legitimate in my sense).[19] It seems to be thought that, even if there is no general obligation to obey a state, we need not worry about the general moral

---

[16] See, for example, Simmons (1979) and Green (2004).
[17] See, for example, Smith (1973), Green (1988), (1996) and (2004), Simmons (1979) and Wolff (1970).
[18] Of course, these questions are not completely separate: an account of the *possibility* of state legitimacy would have to give the conditions under which a state *could* be legitimate. The thought that there *can* be (or are) legitimate states is not so commonly challenged, though, as the thought that states can (or do) have *obligating power*.
[19] Rawls (1993), a particularly influential work that places legitimacy centre stage, is, I think, best read in this way. He takes the existence of the state as a starting point, and so does not seek to show that states' claims to exclusive enforcement rights are justifiable, but rather, given that there are states with exclusive enforcement rights, what conditions they must meet to be acceptable (see footnote 90).

permission to enforce.[20] I do not want to suggest that this is a universal assumption, but just one that seems quite often to be implicitly made. At least, there is a notable lack of concern about legitimacy anarchism, when compared with obligation anarchism. I want to argue that political philosophers have no reason to be so confident about legitimacy, when they are so sceptical about political obligation.

A number of thinkers have argued against 'strong' conceptions of legitimacy, according to which legitimacy is not just the permission to enforce (to which I am stipulatively restricting the term) but also entails political obligation.[21] A good part of the motivation for this seems to be the thought that we do not want to have to say that states generally fail to be legitimate, and, since there are good reasons to be sceptical about political obligation, strong conceptions force us to do so. It seems to be thought that if we do not take obligation to be necessary for legitimacy, we will have no (or less) reason to worry about wholesale legitimacy-scepticism. This is a sensible view to take if we have less reason to worry about the other elements of legitimacy, which are generally taken to include the general permission to enforce. Some of these thinkers provide their own independent arguments for the permission to enforce (which I will come to in chapter 2), but some do not. One place where this assumption seems particularly evident is in William Edmundson's argument for a weak conception of legitimacy (which does not entail obligating power).[22]

Edmundson argues that states can be legitimate in a stronger sense than mine: he argues that what he calls the 'Modest Legitimacy Thesis' (MLT: 'Being a legitimate authority entails that one's authoritative directives create in one's subjects an enforceable duty not to interfere with their forceful administration')[23] can be satisfied. To interfere with the forceful administration of a directive would tend to involve direct disobedience of what he calls an 'administrative prerogative'. An administrative prerogative is an official action involved in the administration of the law targeted at the immediate conduct of an individual on an occasion: for example, the sheriff tells us to be present in court, the judge orders us to pay damages. His argument is that direct disobedience of an 'administrative prerogative' is more presumptuous than simply disobeying a law. He gives a

---

[20] Writers do not always make the distinction between obligating power and legitimacy, but where sceptical doubts are raised, they tend to be about obligating power, not legitimacy.
[21] For instance, Buchanan (2002), Copp (1999), Edmundson (1998) and Greenawalt (1999).
[22] Edmundson (1998).
[23] Ibid. 42

number of reasons for this such as its being 'harder to have an assurance of harmlessness when it is a traffic cop, rather than a traffic law, that one is disobeying'[24]; its being hard or impossible to unconsciously disobey an administrative prerogative (in contrast to a law) and so on. He then also claims that this duty to obey administrative prerogatives is simply the correlative duty of the state's justification right (liberty right/moral permission) to enforce its directives: if one has a justification right to φ, this entails a duty on the part of others not to interfere with your φ-ing and in this case, Edmundson claims, not interfering with the state's enforcing its directives is just equivalent to obeying its administrative prerogatives.[25]

Now, it may seem plausible that direct disobedience of an administrative prerogative issued by an agent of state S is more problematic than disobeying some law made by S *if* S has a right to enforce its directives (and Edmundson's arguments make a strong case for this conditional claim). However, the arguments Edmundson gives seem all to *require* the state to be morally permitted to enforce its laws, commands or directives. If there is no such right, it is hard to see why it would be so presumptuous to disobey an administrative prerogative. And of course, the fact that this duty is the correlative of the right to enforce is of no import if any given state lacks the right to enforce. Edmundson gives no real argument to the effect that states do have such a moral permission. Presumably the thought is that it is, in some sense, reasonably easy or straightforward to establish such a permission. As he says, the most prominent philosophical anarchists, Wolff and Simmons, do not deny the possibility of such a general permission or 'justification right'.[26]

Now, of course, if, as I have said, legitimacy and obligating power are conceptually independent of each other, then a challenge to the latter is not automatically a challenge to the former. So, if we reject the latter we do not yet have any reason to doubt the existence of the former. So why suppose that legitimacy is in any danger? A state's being legitimate in my sense involves it having a general moral permission to do something that seems to be at least prima facie morally problematic, as I will suggest in more detail in chapter 1.

---

[24] Ibid. 50

[25] Ibid. 61

[26] Ibid. 40. Simmons (2001) seems to suggest this kind of openness to states' having such a general permission, although his (1991) appears to suggest the converse. Additionally, Leslie Green (2004), who is, like Wolff and Simmons, a sceptic about political obligation, seems quite clearly to think that, although there are no general and universal obligations to obey the law, at least certain existing legal systems are legitimate (permissibly enforced).

Thus, I think, unless we can find some *argument* to ground the legitimacy of states of a certain kind, we should not suppose that they are legitimate. This thesis will question whether there is such an argument available.

3.   The state

I have said that this thesis is about states and what they are morally permitted to do, but I have not said anything yet about what I take a state to be. I will give a definition of 'state' that I take to be a plausible analysis of our ordinary concept of 'state'.[27] I do not have space, though, to properly defend this definition *as such an analysis*, and it can simply be taken as a stipulative definition if the analytical claim is unconvincing. I propose to understand states as particular kinds of *groups of people.*

It is natural to think about states as *institutions.* If institutions ultimately are groups of people, then this is not in conflict with my proposal. If institutions are groups of people, they will presumably be groups of people structured in particular kinds of ways. States, in the sense I use the word, will be institutions that meet certain conditions (that is, structured groups of people that meet certain conditions). If, on the other hand, institutions are just sets of *structures* themselves, rather than the people organised by these structures, then the state (as I will use the term) will not be the institution(s) but the group of people structured by the relevant institution(s).[28] The state will obviously have to be an agent or group of agents (capable of acting)[29] for it to make sense to talk about what the state may permissibly do, and so cannot simply be *a set of structures*.

I define a state as *a group of people that claims supreme authority and an exclusive right to enforce its commands/directives/laws over a given territory or population, and in fact has extensive enforcement power over this territory or population.* This definition begins with the Weberian idea of the state claiming a monopoly of the legitimate use of force (that is, claiming the right to use force, and also the moral power to determine when use of force by other agents is or is not legitimate).[30] However, as Nozick points out, just *claiming* a monopoly on the legitimate use of force is not

---

[27] There are interesting discussions of this analytical question in Nozick (1974) 22-5 and Green (1988) ch.3.

[28] Copp (1999) characterises the state as an 'animated institution', which is 'an institution or system of offices and roles together with the people who occupy those offices and roles' (6). For him, then, the state is the *combination* of the structures and the group of people.

[29] A group of agents may be capable of acting without constituting a single collective agent; I do not assume that the state is a corporate agent.

[30] Weber (1948) 78

sufficient (and the same will apply to claiming a monopoly on authority).[31] *You* or *I* could *claim* such a monopoly. Doing so would not *give* us any monopoly, nor would it make us a state. It may also not be necessary, since a state could accept the legitimacy of certain unauthorised uses of force without ceasing to be a state (and the US second amendment might be interpreted in this way). I think it is necessary, though, that the state claim the exclusive right to *enforce compliance with its commands/directives/laws*; the state cannot accept the legitimacy of any unauthorised agent doing *this* without ceasing to be a state. I also add that a state must claim supreme authority.[32] This allows us to distinguish groups of bandits from states. I am not certain that groups of bandits do not qualify as states, but including this requirement limits our discussion to the most paradigmatic states.

Finally, Nozick points out that the actual possession of a monopoly on the use of force is not necessary for statehood. This is correct, and nor is it even necessary for a state actually to have the power to enforce *all* of its directives. But what I think *is* necessary is some *reasonably* extensive enforcement power.[33] A group that merely makes certain *claims*, but in fact possesses no actual enforcement power does not seem to qualify as a state (even if these claims are widely accepted by the would-be-state's claimed subjects). This definition, then, has the consequence that this thesis concerns a moral permission that states claim *essentially*. If you find the definition a plausible analysis of the ordinary concept, you will agree that the question of legitimacy is bound up with the very existence of states. (This is not to say, of course, that there could not be illegitimate states: states may make false claims and act impermissibly.)

4. *The generality of the permission to enforce*

I have characterised the question that I am interested in as the question whether states have a general permission to enforce. The interesting question is not simply whether a given state is *ever* permitted to enforce compliance with its directives. We would not, for example, have shown anything very interesting if

---

[31] Nozick (1974) 22. See also Green (1988) 77ff.

[32] The requirement of supremacy does not imply that there can be no limits to state authority (cf. Buchanan (2002) 690).

[33] Nozick also seems to add the requirement that a state must offer protection to all those who live in its territory. I see no reason to suppose that states must be territorial; if a body existed that claimed an exclusive right to enforce over a non-territorially-defined population, why should we not call it a state? I also see no reason to think that a state must offer protection to all those over whom it claims an exclusive right to enforce. It seems to me that a state that didn't do this would still be a state, but a bad state. In any case, whatever is morally problematic about the claim to and exercise of an exclusive enforcement right when made by a body that offers protection to all of its claimed subjects, will presumably also be problematic about the claim to and exercise of an equivalent right by a body that does not offer protection to all.

we showed that it is permissible for the state to push people out of the way of oncoming trains. In cases of impending catastrophe that can be averted only through the use of force, it seems that force is often permitted. Plausibly, this applies to states as well as to individuals. The state's permission to enforce needs to be in some sense *general*. States take themselves, and present themselves, as having a *general* permission to enforce compliance with *all* of their laws and directives. It seems, though, that to demand a carte blanche permission to enforce (a permission to enforce *whatever* directives the state might make, however evil) is to gift the debate to the anarchist unnecessarily. No sensible defender of state enforcement will want to defend the claim that there are states that have a permission to enforce absolutely *any* law or directive they might make. What the legitimacy anarchist wants to deny, it seems, must be somewhere between the claim that a state is *sometimes* permitted to enforce compliance with its directives and the claim that the state is *always* permitted to do so. The former makes life too *hard* for the anarchist, while the latter makes it too *easy*. What, then, is the *generality* of the permission to enforce that the anarchist wants to deny to states?

First, the state's permission to enforce needs to go beyond a permission to enforce compliance with those moral requirements that would exist in a stateless, lawless society, independent of the existence of law and the state (what might be called the *mala in se*). It is not inconceivable that a state might lack even this: it does not seem plausible to think that all of the state-independent moral requirements are permissibly enforceable *in general*, and so it is at least not *obviously* the case that the state has such a permission. However, states claim a lot more than this, and the anarchist would have succeeded in showing something quite significant if she were to show that *all* a state had was a permission to enforce the mala in se (the state-independent moral requirements). To defeat this sort of anarchist, we need to show that the state has some discretion about what it enforces. It needs to have a permission to enforce even some directives compliance with which is not morally required independent of its existence. (This is not necessarily to say that it needs to be permitted to enforce behaviour that is *not* morally required at all; since it could be that certain things become morally required as a result of the state's demanding them.)

Let me suggest, then, this minimal requirement for what would count as a defeat for the anarchist: the state would need *at least* a permission to enforce all of its *just* laws or directives. What I mean by 'just laws or directives' are laws or directives according to which it would be just for the society in question to be structured. I assume that there are more than one possible set of just laws or directives (for a

given society); a state's being limited to only enforcing just laws is consistent with it having discretion about what laws it enforces. In addition, a law's being permitted by justice does *not* mean that what it requires would be morally required *independently* of the law's existence. My suggestion is that, in order to defeat the anarchist, it would be necessary to show that the state has *at least* a permission to enforce all of its just laws. Whether or not it is permissible for any state to enforce some unjust laws is a further question. I will take it that the anarchist has been defeated if it can be shown minimally that a state has a general permission to enforce its *just* laws whatever they may be. If, on the other hand, it can be shown that states lack even this, the anarchist has won.

It could be that it is morally obligatory for the citizens of a certain just state to obey all of its laws and directives. This could be because there is a content-independent and general obligation for all citizens to obey its directives, or it could just be because various other considerations happen to sum together to make obedience to the good and just laws of the state obligatory in all cases (even if it would not be were the laws different, or the circumstances different). If you thought that it is generally permissible to enforce all behaviour that is morally obligatory, then it would be morally permissible for *anyone* to enforce compliance with all of this state's just laws and directives, and so, a fortiori, for the state to do so. Similarly, if you thought that *justice* is always permissibly enforceable, then it would follow straightforwardly that states are permitted to enforce all of their *just* laws and directives.

I think it would be wrong to think that it is generally permissible to enforce all morally required behaviour or that justice is necessarily enforceable (I will elaborate on this below), but setting this aside, I think this scenario would still not really be sufficient for the sort of permission to enforce that a state would need in order to robustly defeat the anarchist. Here, the state's permission to enforce would be no different to that applying to anyone else. States claim not only that their enforcement is legitimate, or permissible, but that *only* their enforcement is legitimate or permissible (or, at least, whatever other enforcement is permissible is to be determined by the state, and this determination can in turn be *enforced* coercively). That is, the state claims an *exclusive* permission to enforce: a permission to enforce all of its laws and to *exclude* others from doing the same.[34] I want to suggest, then, that the state has not fully defeated the legitimacy anarchist unless it can show not only that it is permitted to enforce at least all of

---

[34] On the importance of exclusivity see Wellman (2009) 426.

its just laws, but also that it is permitted to be an exclusive enforcer over some population, that it is permitted to enforce its decisions about who may enforce behaviour of any sort.[35] Thus, let me suggest the following (stipulative) definition of 'legitimacy':

> *Legitimacy.* A state X is legitimate with respect to individual Y if and only if X has a general moral permission to enforce Y's compliance with its commands/directives/laws independent of their content (within some limits, where content-independence can be limited *no more* than to *just* laws) and to enforce its decisions about who may or may not enforce Y's action.

It is this property (and existing states' possession of it) that will be the subject of my thesis.

4.a) Justice and enforcement

Morally obligatory behaviour, I said above, is not always permissibly enforceable. What people are morally required to do is, I think, a separate question from the question what people can permissibly be forced to do. We saw this in the cases of promising and gratitude. Permissible enforcement, then, does not *follow* from moral obligation. There is a popular view, though, that dates back at least as far as Kant, according to which *justice* (a subdomain of moral requirement) is *necessarily* permissibly enforceable, or is necessarily at least a pro tanto ground for enforcement.[36] On this view, it is part of what it is for something to be a requirement of justice that there is at least some pro tanto reason to enforce it. If this were the case, any state would be permitted to enforce all of its *just* laws (or at least it would be pro tanto permissible, unless other countervailing considerations were to block enforcement in certain cases). They are just, and what this means, on this view, is that it is permissible to enforce them. This would not show, as noted above, that all states are legitimate in my sense. If it follows from something's being a requirement of justice that it is permissible to enforce compliance with it, this is to say nothing about *exclusive* enforcement. This view of justice would do nothing to show that any particular institutions are permitted to act as *exclusive* enforcers, that is, to *prevent* others from enforcing the requirements of justice. (For what it is worth, though, I do not find this view of justice at all

---

[35] The state can obviously pass laws permitting its officials to enforce and forbidding others from doing so. Then, *if* these are just laws, and there is an obligation to obey the state's just laws, and if (counterfactually) it is permissible to enforce all obligations, the state will get a permission to be an exclusive enforcer. But the question is what could *make* it permissible for the state to pass such a law. Presumably, if a state is not permitted to be an exclusive enforcer, it will not be permissible for it to pass such laws and they will not count as just laws.

[36] This view finds expression in, for instance, Kant (1999) 6:231-2; Nardin (2004); Valentini (2012a) 597 and Buchanan (2002) 704.

plausible: if requirements of justice are roughly those things that people owe to each other, there is no obvious reason to think this has anything to do with what they can legitimately be forced to do. There seem to be many requirements of justice, such as keeping promises, that do not give even a pro tanto ground for enforcement.)

5. The particularity requirement

Simmons, in his investigation of political obligation, introduces what he calls the 'particularity requirement'. That is, he says that he is 'only interested in those moral requirements which bind an individual to one *particular* political community, set of political institutions, etc'.[37] Thus, a general duty to support just institutions, for instance, would not provide the sort of political obligation that he is concerned with. Such a duty applies just as much to other just states as to one's own; it provides no *special* bond between you and *your* state. If this were the only sort of political obligation we had, then a state's being *your* state would add nothing to the moral calculus (though of course, its geographical proximity might). I think that a similar requirement is just as important, if not more so, for *legitimacy* in my sense. If what the defender of legitimacy wants to show is that at least some of our actual, existing states have the kind of general permission to enforce that they are standardly taken to have (and not just that it is *possible* for a state to be legitimate), it will be necessary to be able to show for at least some actual state that *that* state has a general permission to enforce. It will not do simply to show that it is permissible for *some* state to enforce its directives. All that would show is that it is *possible* for a state to be legitimate. Legitimacy is something that must be shown state by state: for each state where the question arises, it must be shown that *that particular* state is permitted to enforce.

6. Outline

We now have an idea of what is at stake for the legitimacy anarchist. The legitimacy anarchist denies that any existing states are legitimate with respect to their subjects in the sense just given above. The aim of this thesis, as suggested above, is to argue that this anarchist view should be taken seriously in a way that it often is not. What I hope to do, then, is to cast some doubt on the assumption that at least some existing states are legitimate in this sense. In particular, the focus of this thesis will be on the role of feasibility considerations in arguments purporting to establish state legitimacy. I will argue that all of the arguments that

---

[37] Simmons (1979) 31

could plausibly establish legitimacy for some existing states depend on premises about feasibility. However, the nature of feasibility is such that this dependence is enough to shed some doubt on the success of these arguments: given the account of feasibility I will give, the truth of the feasibility premises that would make the argument go through should not be taken for granted.

In the first part of the thesis, I argue for what is also an interesting conclusion in its own right: state legitimacy cannot be established a priori (independent of empirical considerations) but rather depends on certain feasibility facts. I proceed by first arguing that a plausible line of thought that might underlie the belief that legitimacy is on safe ground only works with the inclusion of feasibility premises. There are some prominent arguments, however, that the legitimacy of a certain kind of state can be established a priori (and so a fortiori without feasibility premises). I thus turn to showing that these arguments fail.

In the second part of the thesis I offer a multivocal account of feasibility, according to which there are many possible sharpenings of feasibility (ways of making the concept precise), no single one of which is obviously privileged over the others (for the purposes of political or moral theory or more generally). Since the feasibility premises required to establish legitimacy will not come out as true on *all* such sharpenings, some argument will be needed that the sharpenings that would allow the arguments in question to go through *are* ones on which the premises are true. Further, I argue that theory constrained only by *unrealistic* feasibility constraints can be worthwhile, and thus that the conclusion that state legitimacy cannot be established a priori is itself important and useful.

Chapter 1

In the first preliminary chapter I give an account of what I mean by 'enforcement' and I argue that a common intuition gives us reason to think that there is some moral presumption that would need to be overcome to establish legitimacy in the above sense.

Chapter 2

One plausible line of thought that might underlie the assumption that state legitimacy (for at least some actual states) is on safe ground is that it follows from a property that at least some states possess reasonably uncontroversially, namely justification. I consider the most obvious interpretations of what might be meant by the claim that a state is justified, and argue that all of the properties identified are such that we are only warranted in assuming *both* that some states possess the property *and* that legitimacy follows from the property, if certain feasibility

premises hold. I also argue that most of the explicit arguments for state legitimacy have this underlying form (and that those that do not are unlikely to establish the legitimacy of existing states).

Chapter 3

In chapters 3 and 4 I address some arguments that claim to show that state legitimacy (for states that meet certain conditions) can be established *a priori* (independent of any empirical considerations, and so independent of any feasibility considerations). First, I address one interpretation of the Kantian argument in the *Doctrine of Right*, that given by Arthur Ripstein.[38] On this interpretation, what makes exclusive state enforcement necessary for a condition of right (and so permitted by right) is the need for *assurance* for there to be rights. I argue that this argument cannot show that exclusive state enforcement is necessary for assurance without the aid of feasibility premises.

Chapter 4

I then turn to another a priori argument for the legitimacy of a certain kind of state, one version of which is given as an alternative interpretation of Kant by Japa Pallikkathayil, and another version of which is given by Philip Pettit.[39] On this argument, exclusive state enforcement is required for people's freedom as independence (or non-domination) in society (that is, people's not being *dependent on* or *dominated by* the will of others). I argue again that this cannot be shown without feasibility premises. Thus, the first part of the thesis concludes that state legitimacy cannot be established a priori and that plausible lines of argument that might ground it depend on feasibility premises.

Chapter 5

In this chapter I offer a general account of the concept of feasibility. The account I offer is *multivocal* and *possibility-based*. There is a whole range of possible (binary) sharpenings of the term 'feasible', each of which selects a range of facts of the world to hold fixed (and feasibility is defined in terms of possibility given this range of facts). No single one of these possible sharpenings, though, is obviously privileged as giving the appropriate understanding of 'feasibility' *tout court*.

Chapter 6

I argue that the upshot of this account of feasibility is that the kind of arguments for state legitimacy that make use of feasibility premises only go through if these feasibility premises are true on a sharpening of feasibility that also licences

---

[38] Ripstein (2009) and Kant (1999).
[39] Pallikkathayil (2017) and Pettit (2012).

something like an "'ought' implies 'feasible'" principle (i.e., a sharpening of feasibility that is a constraint on the truth of claims about what we ought to do). Some argument would be needed to show that this is the case, so we should not simply *assume* that legitimacy is on safe ground. Further, I argue that the conclusion that states *lack* a general and exclusive permission to enforce given certain *unrealistic* sharpenings of feasibility (those on which the necessary feasibility premises are not true) is one that it can be *useful* to learn (for the purpose of action guidance). In particular, we learn something about how we ought to treat and think about state enforcement.

Chapter 1:
Preliminaries

A. What is enforcement?

I have defined 'legitimacy' as a state's general and exclusive moral permission to enforce compliance with its directives. I have not said, though, much about what I mean by 'enforcement'. The rough idea is intuitive. It is one thing to have the right to issue directives or commands and to have the ability to make directives or commands that others have a reason, or a duty, to follow. It is quite another thing to be permitted to *make* those others comply with one's directives or commands, to make things such that they have no choice but to comply. For the sake of clarity, though, it is worth giving a more precise definition of 'enforcement'. My proposed definition of enforcement is best introduced with the aid of Serena Olsaretti's distinction between voluntariness and freedom.[40] She says that 'a choice is voluntary if and only if it is not made *because* there is no acceptable alternative to it'. Freedom, on the other hand, is about the options we face. She gives two examples to illustrate this distinction:

> *The Desert City.* Daisy is the inhabitant of a city, located in the middle of a desert, which she is free to leave. However, Daisy, who would wish to leave, knows with absolute certainty that if she leaves the city, she will not be able to survive the hardship of the desert and she will die. Her choice to remain in the city is not a voluntary one.

> *The Wired City.* Wendy is the inhabitant of a city fenced with electrifying wire, which she is unfree to leave. However, her city has all that anyone could ever ask for, and Wendy, who is perfectly happy with her life there, has no wish of leaving it. She voluntarily remains in her city.[41]

Coercion, Olsaretti notes, need not reduce freedom (if coercion is understood as pressure on the will by means of a threat or similar), as in the case of a bluff threat that successfully coerces A into not doing x, but where, in fact, doing x would have no adverse consequences for A. Coercion (in this sense) does, though, always undermine voluntariness. Coercion, for Olsaretti, is a form of *forcing*, where forcing is defined along these lines:

> *Forcing.* A is forced to φ if and only if A does φ involuntarily.

---

[40] Olsaretti (2004) pp. 138-50
[41] Ibid. 138

> B forces A to φ if and only if B makes things such that A does φ involuntarily.

Coercion, though, she says, is only one form of forcing, one way in which people's choices may be made involuntary. She argues that there is no reason to fetishize coercion: a concern with voluntariness should also justify taking seriously cases of non-coercive non-voluntariness. Coercion (in the above sense) is a *technique* for forcing. (It is for this reason that I talk about enforcement rather than coercion. On what Scott Anderson has called the 'pressure account', according to which it involves the putting of pressure on the will of another using threats or similar, it is only one form of a broader single kind of activity, of which (for instance) the use of physical force is also a variety. This seems also to be true on Anderson's own 'power approach', according to which coercion involves 'a significant disparity in power'.[42] I do not have space to go into this in any detail, but there seem to be instances of the same kind of phenomenon where there is no non-trivial disparity in power.)

I suggest, then, that 'enforcement' can be defined thus:

> *Enforcement.* Agent A enforces Y's doing φ (being P) if and only if A intentionally makes it the case that either a) Y does φ (is P) involuntarily or b) Y does φ (is P) voluntarily but if Y had not done φ (been P) voluntarily Y would have done φ (been P) involuntarily.

Enforcement, as understood here, involves either actual or counterfactual instances of forcing. A's doing φ is enforced just if, were A not to φ voluntarily, A would be forced to φ. In both of Olsaretti's cities, if somebody is responsible for *making it the case* that Daisy and Wendy cannot leave their cities, they will have enforced Daisy's and Wendy's remaining in the city. Daisy is *forced* to stay in her city; her choice is not a voluntary one. Wendy is not forced to stay, since her choice *is* voluntary, but if she had not stayed voluntarily, she would have been forced to stay. Voluntariness for Olsaretti, as I have noted above, requires that an action *not* be done *because* there are no acceptable alternatives. This, of course, leaves some important questions unanswered. For one thing, it necessitates some account of what counts as an *acceptable* alternative. For Olsaretti, the standard is an 'objective standard of well-being'. I will not attempt to give any account of acceptability, so my definition of 'enforcement' is by no means complete. I will just assume (as is intuitive) that there are things that others can do to make certain alternatives otherwise open to us *ineligible* without making those alternatives

---

[42] Anderson (2010) 1

strictly unavailable.[43] This is not supposed to be an account of our ordinary language concept of enforcement, but rather a definition of a technical term. The phenomenon picked out here is the one with which I will be concerned.

There is what might seem to be a problem for this definition of enforcement. Since my definition of 'enforcement' makes it a success term (an agent's doing φ or being P is only enforced if they *do* do φ or are P), it doesn't seem to allow directives to be enforced on those who disobey. If you disobey a directive demanding you φ, nobody has enforced your φ-ing, even if they then punish you for failing to φ; at best, they have *failed* to enforce your φ-ing. It might be thought, though, that punishment for disobedience is a paradigm case of state enforcement, so there is something wrong with the definition. I do not think that there is a problem here. Certainly, punishment for disobedience involves enforcement, but my definition need not deny this. I think that cases in which things are made such that the subject of a directive *does* do what is demanded, whether they like it or not, are the basic cases of enforcement. What my definition must say is that punishment for violation of a directive does not strictly speaking enforce *that* directive, even though it might be natural to say that it does. It is right that cases of punishment for disobedience are paradigm cases of a state's enforcement of its directives. However, what is being enforced in such a case is the part of a directive that says something like 'if A fails to φ, A will be made to do/be, prevented from doing/being ...'. For example, when A is imprisoned for murder the state *does* enforce a directive on A: it makes it the case that A is in prison for some term, irrespective of whether A wants to be. What it does *not* do is enforce A's not murdering. The directive that says, 'do not murder' is only successfully enforced on those who in fact do not murder and here it is by the *threat* of imprisonment that it does this.[44]

1. *Enforcement is not necessarily moralised*

My definition of enforcement is non-moralised. The moral wrongness of an act of enforcement is not a conceptual necessity (nor even is its being prima facie morally problematic). It also seems not to be a necessity of any other sort. Enforcement as I have defined it is just too banal a phenomenon to be morally

---

[43] It may well be that *eligibility* or *acceptability* comes in degrees. If this is so, it will make enforcement a vague term. A will only enforce B's φ-ing if A makes it the case that if B were not to φ voluntarily, she would φ because of the absence of *sufficiently* eligible alternatives.

[44] Similarly, if I make it very costly for you to φ, but not ineligible, I do not enforce your φ-ing, but I may well thereby be enforcing something on you: I enforce your not doing certain combinations of actions (φ-ing without paying the cost).

problematic *as such*. People have things enforced on them all the time, at least if we want to allow (as we do) that one's not doing something, or one's being something, can be enforced. For instance, I have no acceptable alternative to not going to the moon, not speaking to Boris Johnson over breakfast and so on. Thus, if I fail to do these things, as I will, we can say that my not doing these things has been enforced. This cannot be morally problematic. Perhaps an agent's intentionally enforcing things on others is what raises moral problems. However, even this might seem to be too banal. I might intentionally arrange things such that you have no acceptable alternative to not talking to me today. I can do this quite easily by avoiding being in the same place as you. Is there something morally problematic about this? It does seem plausible to think that there is not.

I do not claim, then, that my concept of enforcement is the *only* useful concept in the area. If there is a similar phenomenon that *is* necessarily morally wrong or problematic, then it will certainly be useful to have a corresponding concept. I am not sure, though, what this concept would be. It is not clear that standard 'pressure' accounts of coercion identify phenomena that are always morally problematic.[45] But even if they do, it seems like they will miss out cases of enforcement that are morally problematic in the same way. (It seems plausible that the use of direct force can be problematic in the same way as pressure on the will.) Thus, identifying the fundamental phenomenon that is necessarily morally problematic is another task. Given, then, that we have not identified an activity that is always morally problematic, the state's enforcing a directive it makes is not something that must always defeat a presumption of impermissibility. It is not the case that enforcement always automatically stands in need of a justification. However, as I will now explain, I think there *is* a presumption that needs to be defeated in order to show that a state (or any agent or agency) has a *general* permission to enforce compliance with its directives of the sort I have described.

B. The presumptive moral complaint

I find it intuitively hard to doubt that the kind of general and exclusive permission to enforce that I have described involves a permission to do something that would not ordinarily be permissible for individuals. I have little intuitive doubt that there is some pro tanto moral wrong or moral presumption (even though it is difficult to say *exactly* what it is) that must be overcome for a state (or any other agent or institution for that matter) to have this kind of general

---

[45] Anderson's (2010) 'power approach' to coercion, mentioned above, does not seem to either.

permission to enforce. I am here merely reporting my own personal intuition, but I do think that this feeling is quite common. The project of my thesis is motivated by this assumption that the state's general permission to enforce stands in need of justification.

The assumption is that there is something morally troubling about at least some of the actions of a state that would be permitted by a relevantly general permission to enforce, and that this necessitates some *argument* to show that a state can or does have such a permission. If we make such an assumption, then the failure to find such an argument would constitute good reason to doubt that states have a general permission to enforce. If the reader does not share this intuition, they need not be so troubled by such a failure or by the arguments of this thesis. They may accept, as this thesis argues, that there is good reason to doubt whether there is a successful argument available for the general permission to enforce for any existing states, but yet think this of little interest since no such argument is needed. I take it, though, that those who have advanced arguments for the possibility of a state's general permission to enforce *do* share my intuition. They consider it necessary to give such an argument, I presume, because they think that the state is permitted to do something that stands in need of justification, something which is prima facie problematic, and so which needs to be defended by argument. I do not think, then, that in making such an assumption, I am assuming anything very controversial.

The intuition that I think is widely shared is not that the general permission to enforce involves a permission to do something that is *wrong* (that would be an intuition that there can be no general permission to enforce), nor that it is hard to establish. The intuition is merely that there is something pro tanto morally problematic about at least a part of what such a general permission would permit. One might share this intuition but think that the pro tanto moral objection is easily overcome by strong countervailing considerations. I must also make clear that the intuition is not that enforcement, or state enforcement, is necessarily pro tanto problematic. On the contrary, I noted in the previous section, it seems that enforcement *as such* is not problematic. And I do not suggest that there is something necessarily different about enforcement when carried out by *states*.

Rather, the intuition is that there is a pro tanto complaint against *certain sorts* of enforcement that a state would need to be permitted to carry out in order for its permission to enforce to be sufficiently *general* (for it to be permitted to enforce at least *any* of its just directives whatever they may be). At least some possible cases

of enforcement of just directives are morally problematic. Some of the kinds of things that a state might enforce in enforcing compliance with its just directives are pro tanto impermissible to enforce. (In fact, my intuition is that *most*, or at least *many*, of these things are pro tanto problematic to enforce, but this is more than is necessary.) A state could have a *general* permission to enforce without being permitted to use any *method* of enforcement, so the idea is not just that some possible *instances* of enforcement of just directives are presumptively wrong, but rather that some possible just *directives* are presumptively wrong to enforce. There is something troubling (not wrong, but in need of justification) about at least some subset of states' enforcement, and there is still something troubling about a subset of states' enforcement even if the only directives they enforce are just ones. This is not a problem with enforcement *as such*, but with some significant subset of the enforcement that states do. A general permission to enforce is a permission to do things that must *include* the things that are troubling. Thus, in order to defend it, this intuitive moral presumption will need to be overcome.

Although I think I am not alone in feeling the intuitive compulsion to think that there is a presumptive moral complaint that states' general permission to enforce must overcome, it is not easy to identify exactly *what* this complaint is. Niko Kolodny, in a fascinating recent piece, has raised this question.[46] He suggests (rightly, I think) that it is widely thought that there is a general pro tanto complaint against 'relations of rule' (which, in his use, include the enforcing of commands over a group of subjects, and the claim of an exclusive right to do so), but asks what this could amount to.[47] Though it is an interesting and important question, I will not attempt to answer it in this thesis. Instead, I will rely on the intuition that there is *some* such complaint. The difficulty of identifying *what* it is, though, might lead one to wonder whether the common intuition is simply a mistake. This appears to be the case for Kolodny himself. (He identifies a possible candidate complaint that he thinks most likely to do the job, but expresses doubt as to whether there is any complaint of the sort he is looking for at all.) His reasons for uncertainty stem from his argument that what seem the most obvious, or commonly identified, candidate complaints do not stand up.

If there in fact is no presumptive complaint against states' generalised enforcement, the general and exclusive permission to enforce will not stand in need of argument, and so the success or failure of such arguments will be of little

---

[46] Kolodny (2016).
[47] Ibid. 35

importance. For this reason, I think it worth taking a brief detour to look at a handful of possible candidate complaints and to suggest that there are a few that have some intuitive plausibility and which we have no immediately obvious reason to reject. Since in (at least) all but one case, Kolodny denies this, I will argue that he does not convincingly show that these candidates could not provide the sort of pro tanto complaint needed. My aim is just to suggest that there is reason to think that grounds for the common intuition are out there to be found. I believe common opinion is on my side, but I think it is worth saying a little more to dispel doubts about this, since Kolodny has offered a powerful challenge to the standard belief. I do not aim, though, to make a serious attempt to identify the grounds for the common intuition in any very satisfying way. I hope merely to make my assumption seem like an acceptable one to make.

1. *Deontological complaint*

    i) *Against forcing*

One possibility Kolodny considers is that there is some sort of deontological constraint that is violated by state relations of rule.[48] There are certain things that we may not do to a person even for the sake of producing a greater good. What might the deontological constraint be? It could be that there is a constraint against *forcing* of some kind. Forcing, as I defined it (following Olsaretti), is making another do something against their will, involuntarily. It is quite intuitively tempting to think that there is some sort of pro tanto deontological rule against making others do things against their will. However, just as we saw that enforcement is not necessarily morally problematic, it may be that forcing is not either. If I sit in a chair that you wanted to sit in, I make it the case that, against your will, you do not sit in the chair. It is natural to think that there is nothing morally problematic about my doing this, and so, if this counts as forcing (and it seems to, given the definition), it is possible for forcing to be unproblematic.

However, even if there is no deontological constraint against forcing *as such*, it is plausible to think that there is *some* deontological constraint against a wide range of forcings. It is natural to think that in more paradigmatic cases of forcing, such as where whenever you attempt to sit on your favourite chair I deliberately wrestle you away, there *is* some pro tanto deontological constraint that is being violated. It is difficult to say what exactly the principled distinction between these kinds of cases is, but that this distinction is difficult to identify does not mean that

---

[48] Ibid. 45-60

there is *no* such distinction. If there is a deontological constraint of this sort, then the kind of general permission to enforce that we are looking for would have to overcome it. A general permission to enforce must be a permission to make it the case that the state's subjects will comply *whether or not they want to*, so it will need to include a permission to force them to comply when they do not do so willingly. If there are unwilling subjects, it certainly is not the case that the only sort of forcing the state would have to engage in to enforce compliance with its just directives would be like sitting in another person's desired chair. If there is a deontological constraint against some subset of forcings (and it is quite intuitively natural to think that there must be), then it will surely be a presumption that the general permission to enforce must defeat.

ii)   *Against methods of forcing*

It could be that the only way of drawing a distinction between forcing that is morally problematic and that which is not is in terms of the *method* used. Kolodny, however, argues that there is no constraint against a method of forcing that is necessary to a state's enforcing its just directives. I will argue that he fails to show this and so a deontological constraint against some necessary aspect of the methods used for state enforcement remains a possible source for the presumptive moral complaint.

One plausible thought is that there is a deontological constraint against the use of *violence* or *physical force*. Kolodny suggests a 'Force Constraint' as a candidate complaint against state enforcement:

> *Force Constraint:* It is impermissible to use force on someone as a means to, or foreseeable side-effect of a means to, a greater good.[49]

If the general permission to enforce requires a state to be permitted to use force or violence, then this constraint will establish a presumption needing to be overcome. However, a state's having a general permission to enforce does not require it to be permitted to use *any possible method* to enforce its just directives, so if it is possible for a state to enforce all of its just directives *without* the use of force or violence, then this constraint will not be what we are looking for. Kolodny points out that it *is* possible for a state to enforce its directives without the use of force. He describes a society where this happens, the 'Omittite Empire':

> Their emperor, the Guardian of the Ladder, does not put violators of his directives in prison or build prisons around them. He doesn't need to. This is because each Omittite, to survive the elements, must descend into his naturally

---

[49] Ibid. 45

carved hole each night. Every morning, the Guardian drops the ladder into each hole to enable its occupant to climb back up. His deterrent is simply to withhold the ladder, confining the occupant there for a fixed period. Suppose an Omittite, "Holton", violates some directive, and so the Guardian, as announced, does not drop the ladder into Holton's hole for several months. This isn't a use of force or an "active harming", it is simply a failure to aid.[50]

However, Kolodny says, it is quite plausible that there is also another deontological constraint against refusing to aid plus something like the following:

> *Non-Aid Constraint*: If one is otherwise required to aid someone, it is not sufficient to release one from this requirement that by refusing to aid that person, one can use or affect that person as a means to a greater good.[51]

Kolodny argues, though, that the Guardian's enforcing deterrents need not involve using people as means in this way (and so this is not a constraint that the general permission to enforce must overcome). If the aim of the Guardian was to make Holton (a violator of some directive) suffer in order to show others that the Guardian's withholding of the ladder is a bad thing, it would be a refusal to aid for the sake of using Holton as a means. However, the point of the deterrent is not to do this (we can assume that it is obvious to inhabitants of this society that withholding the ladder is a bad thing), but rather to make it clear that the Guardian's threat is credible. Thus, Kolodny says, 'nothing that happens to Holton as a result is part of the Guardian's means to the greater good'.[52] Holton's *suffering* is not caused by the Guardian as a means to a greater good. (The Guardian could do without it: if confinement were in fact a benefit to Holton, the Guardian's deterrent aim would not be set back.)

Presumably, Kolodny's thought is that in treating Holton as above, one does not refuse to aid Holton for the sake of *using or affecting him* as a means to a greater good. What actually happens to Holton is not important to the Guardian's aims. It is not clear, though, what would be wrong with a simpler principle:

> *Non-Aid Constraint'*: If one is otherwise required to aid someone, it is not sufficient to release one from this requirement that *the refusal to aid* that person can be used as a means to a greater good.

This principle would give Holton a complaint against the above treatment.

The reason Kolodny suggests the first principle is that he thinks it explains why 'we may refuse to give life-saving medications to the one in order to have it to give to the five. But we may not refuse to give life-saving medication to the one in

[50] Ibid. 46
[51] Ibid. 46
[52] Ibid. 47

order to learn from the progress of his disease how to save the five from it'.[53] But *Non-Aid Constraint'* seems capable of explaining this too. In refusing to give life-saving medication to the one in order to give it to the five, we do not use *the refusal to aid the one* as a means to a greater good. Our *refusing to aid the one* is not involved in our aim, it is not an essential part of our plan that *we refuse aid to the one*, nor something that we desire. If it were possible to aid the one without giving her some of the medicine that we need to save the five, we might well do so (it would be consistent with our aims). Thus, though perhaps we might be said to aim at refusing *medicine* to the one (since, given the stock of medicine, it *is* essential to our plan, to saving the greatest number, that we refuse it to the one), we cannot be said to aim at *refusing aid* to the one. In this sense, then, the refusal to aid the one in the first case is something like a foreseen but unintended consequence of keeping the medicine for the five.[54]

On the other hand, in refusing to give medicine to the one in order to learn from the progress of his disease, the refusal to aid *is* a necessary part of our plan; if we desire to learn from the one's disease, we must desire not only that we not give our medicine to the one, but also *that we not aid* the one. It is essential to our achieving our aim that we not aid the one, whatever form that aid might take. Here, then, unlike the former case, we aim directly at *not aiding* the one as a means to a greater good. The Guardian's refusal to drop the ladder for Holton is like the latter case, not the former. It is not just the refusal to drop the ladder that is necessary to the Guardian's plan, but also *the refusal to aid*. If it were possible for the Guardian to aid Holton to leave his hole without giving him the ladder, the Guardian could not do so without defeating the object of refusing the ladder. Though Holton's *suffering* as a result of being refused aid is not important to the Guardian's aims, his being *refused aid* is. It is this, his refusing aid to Holton, that makes the Guardian's threats credible. Though the Guardian's method of enforcement does not violate *Non-Aid Constraint*, then, it *does* violate *Non-Aid Constraint'*. Thus, contra Kolodny, there seems to be a not implausible deontological complaint that is violated even in the case of the Omittites.[55]

---

[53] Ibid. 46

[54] See Foot (2002) 20 and 24

[55] Kolodny goes on to argue that the deontological complaint against the use of force can be met by state enforcement, since, roughly, it is best understood as being lifted when the target has had adequate opportunity to avoid the use of force. Adequate opportunity, he says, is determined by the costs of requiring consent in different kinds of cases. I do not have space to discuss this argument properly, but I do not think it is persuasive. It is quite tempting to think that there is a deontological constraint against the use of force that, in

Kolodny fails to show, then, that a deontological constraint (or combination of deontological constraints) could not be the source of the complaint against state enforcement. It seems that making people do things involuntarily will require either the use of physical violence (or the active imposition of harms, if this does not count under the former), or the refusal to aid, or the threat of one of these things. And in order for the threats to function, it will be necessary actually to impose deterrent harms (either by violence or by refusal to aid). Thus, given that a general permission to enforce needs to include a permission to enforce against subjects who are not willing to comply, it needs to include a permission either to use physical violence or to refuse aid in some circumstances. And it is not implausible to think that there are deontological constraints against these things. I have not shown that there *are* such constraints, but it remains, I think, an open possibility.

2. Non-deontological complaint against loss of freedom

If you enforce compliance with a command or directive, this necessarily involves removing or making ineligible certain options for the person whose compliance is enforced. There is a natural conception (or several natural conceptions) of freedom according to which the removal of options is a paradigm case of the restriction of freedom. A hard-line such conception says that *only* removing options reduces a person's freedom. On such a view, so long as an option is strictly open to an agent, no matter how high the cost, it is something they are free to do. However, it is quite natural to think that making an action ineligible, or unacceptable, can also be sufficient to remove an agent's freedom to perform that action. It is odd to say, for example, that the victim of a highwayman, threatened with death if they refuse to hand over their money, is *free* to refuse. There is some sense in which they are, but there also seems to be a very natural sense in which they are not. In this sense, then, enforcement will always restrict freedom.

If we think that freedom of this sort is something that has some value, then an agent's or an institution's having an extensive capacity to remove or make ineligible options of others will be a source of disvalue. I think that this could be enough to show that there is a presumption against the permissibility of exercising such a capacity. Something's simply leading to an overall loss of value all things considered is not enough to establish a presumption against it (unless

some cases at least, is independent of the interest that others in society have in not allowing an individual the opportunity to avoid the use of force.

we are consequentialists). However, I think a state's enforcement leading to a *significant* disvalue may be sufficient. A state's enforcing a wide range of directives covering various domains of life seems likely to cause a *significant* loss of freedom. If freedom is of sufficient importance, this may be enough to establish the sort of presumption we are looking for. (Although Kolodny wants a complaint that persists even for a state that achieves the best possible distribution of goods, which would have to be one where any loss in freedom is outweighed, we do not need to follow him in this. If such an ideal state is able to defeat the presumption, this does not show that there is no such presumption, nor that other less ideal states would be able to overcome it.)

3. Subordination or domination

Kolodny suggests that if there is a problem posed by the relations of rule involved in the state at all, it will be a problem to do with subordination. The challenge the state has to meet, he suggests, is to show that it is compatible with social equality, where this is understood as 'not being subordinated to any individual as an inferior to a superior'.[56] Subordination for Kolodny involves being subjected to the greater power and de facto authority of another individual.[57] It is an ideal of equality since it is violated when there are asymmetries in power and de facto authority between individuals, but it is a *relational* ideal of equality since it concerns equality in how individuals relate to one another, not equality in distribution of some good. Because subordination understood in this way involves social equality of this sort, it is not a problem for states' generalised enforcement as such. The complaint would not arise for an ideal perfectly democratic state, one where everyone has perfectly equal influence over the state's enforcement decisions. Where everyone influences the state in exactly the same way, there is no inequality of power or de facto authority. If relations between individuals are perfectly symmetrical, there can be no subordination understood in this way. Thus, Kolodny's suggested complaint is not one that would need to be defeated to show that it is *possible* for a state to have a general permission to enforce, since it is not a complaint that applies to an ideal perfectly democratic state.

However, there might be another sort of complaint against state enforcement to do with subordination or domination, which *does* apply to state enforcement as such. There does seem to be something troubling, or in need of justification, about another agent having control over a wide range of what one will or will not do,

---

[56] Ibid. 68
[57] Ibid. 63. See also Kolodny (2014b).

when this power is unequalled by your own power over that agent. For a state to have a general permission to enforce, it will need to be permitted to have precisely this sort of widespread control over what its subjects will and will not do. Whatever exactly it is that is troubling about such an asymmetrical power, it is not clear that it is any less troubling when the agent that has control over you is a *group* agent. If the only complaint is about inequality between individuals, then there is no problem in being controlled by a group if there is no inequality between you and any of the individual members of the group. But it is plausible to think that that is *not* the only complaint, that there is something wrong even with being controlled in this way by *groups*. And further, if a group agent's having such control over you is problematic, there again seems no reason to think that it should make a difference that one is a member of the group that has control (unless you have the same sort of control over the group agent that it has over you).

I am not sure what exactly this complaint is, if there is one. It could be that it is a complaint against domination, as understood by republican thinkers, or against dependence on the will of another, as in Kantian conceptions of freedom.[58] Domination is usually understood as something like the uncontrolled power to interfere in the choices of another, while being dependent on the will of another is having one's ability to pursue one's own purposes be subject to the choices or ends of another. These are both introduced, in the republican literature, and in Kant, as conceptions of freedom, so it could be that the subordination or domination complaint is another form of complaint against loss of freedom. A state's being able to enforce its just directives whatever they may be involves it having a discretionary ability to interfere in the choices of its subjects to enforce the directives it chooses. This seems to involve both domination and dependence on the will of another. There is no obvious reason why one could not be dominated by, or dependent on the will of, a group agent. I am not sure that these are the best ways of understanding the complaint, but they are popular ideas, and if they are genuine moral concerns, they could explain the moral presumption against states' general permission to enforce. Alternatively, there could be some other way of fleshing out the trouble with being subject to the control of another agent.

---

[58] On the former, see for instance Pettit (1996), (1997) and (2012). On the latter, see Kant (1999) and Ripstein (2009).

Conclusion

I have sketched very briefly some rough kinds of grounds that there could be for the intuition that there is a moral presumption that a general permission to enforce would need to overcome. The purpose of this was not to provide an account of what the true grounds of this intuition are, but just to make it seem plausible that there are some grounds to be found. One might still think that a state's general permission to enforce stands in no need of justification at all. If this is your view, you will find nothing to trouble you in my thesis. I hope, though, to have made it not seem foolhardy to set aside this view and to work from the assumption that we should not assume that states have a general and exclusive permission to enforce compliance with their directives unless we can find some justification for it.

One final thing that is worth noting before we move on is that none of the suggested complaints seem obviously to go away if the agent doing the enforcing is a group agent, and if the 'enforcee' is a member of that group. First, if there is a deontological constraint against certain kinds of forcing, there is no immediate reason to think this will not apply equally to group agents (if such things exist). If there are group agents, they are agents and so will presumably be subject to roughly the same range of deontological constraints as other agents. Further, nothing seems to change when the 'enforcee' is a member of the group doing the enforcing, even a member with equal voting rights in determining what the group will do. A group agent is capable of forcing other agents, including its own members to do things, and if there is a deontological constraint against certain kinds of forcing, there is no reason to suppose this will not apply to groups forcing their own members to do things. Second, if the loss of eligible options involved in enforcement amounts to a problematic sort of unfreedom, this is plausibly no less the case when caused by group agents, or group agents of which one is a member. Finally, we have already noted that, whatever exactly the subordination complaint is, there is no obvious reason to suppose that it does not apply equally when the subordinating agent is a group agent. Further, it seems similarly plausible that if one can be subordinated by a group, being a *member* of the group does not rule out the possibility. Being a member of a group does not rule out that group having extensive unreciprocated control over your life.

It seems quite plausible also that, whatever exactly *is* problematic about state enforcement, it will not make any difference that one is an *equal* member of the state. It is perfectly possible for group agents to act impermissibly towards their

members, even when their members have an equal say in the group's decision of how to act. (To hold the contrary would be to hold quite a radical, and implausible, view about the virtues of majoritarian decision making.) Roughly this point is made about democracy by John Stuart Mill at the start of *On Liberty*:

> The will of the people ... practically means the will of the most numerous or the most active *part* of the people; the majority, or those who succeed in making themselves accepted as the majority; the people, consequently, *may* desire to oppress a part of their number; and precautions are as much needed against this as against any other abuse of power.[59]

There may be specific cases where one's being a member of a group that treats you in some way is relevant to showing that the group does not wrong you. But it does not generally rule out the group's being able to wrong you, and there seems no reason to *assume* that it does in the case of state enforcement. If there is a presumptive complaint against some of the kinds of enforcement that a general permission to enforce must include, there is no immediate reason to suppose that this complaint goes away simply because those whose action is enforced are equal voting members of the group agent doing the enforcing.

---

[59] Mill (2015) 7-8

Chapter 2:
Feasibility and arguments for legitimacy

A. The implicit argument from justification

I said in the introduction that there is a somewhat widespread assumption that state legitimacy (by which I mean, roughly, states' general and exclusive moral permission to enforce compliance with their directives) is on safe ground, that is, can be taken for granted (for at least some existing states). It could be that this assumption is explained by the belief that there is no moral presumption that state legitimacy would have to overcome. If this is the case, my thesis will not have much to say about it. I doubt, though, whether this is the explanation. I suspect that state legitimacy is taken for granted by many who would *not* think that there is nothing even pro tanto morally problematic about states' exclusive and generalised enforcement. I suspect that the legitimacy of some 'decent' existing states is often taken for granted rather because it is thought that these decent states succeed in overcoming the moral presumption.

In this chapter I will address a line of thought that I think often underlies this implicit assumption. It also seems to underlie one of the primary families of *explicit* arguments for state legitimacy (the family of contractualist and necessity-based arguments). It is a line of thought that has some plausibility and seems as likely as any to warrant taking state legitimacy for granted. My claim will be that if this line of thought *can* warrant the belief that some actual states are legitimate, its doing so is dependent on certain premises about feasibility. I will then argue that another prominent sort of argument (consent-based arguments) for state legitimacy also cannot succeed without the aid of feasibility premises. Thus, the conclusion of this chapter will be that the most plausible arguments for state legitimacy depend for their success on feasibility premises. However, there will remain one group of theorists who think that the legitimacy of a certain sort of state can be established *a priori* (and since feasibility premises are *empirical*, this means without feasibility premises). These arguments will be the subject of the next two chapters.

The thought that I think underlies the common assumption that the legitimacy of some states is on safe ground is the thought that legitimacy follows from a property that some states are considered relatively uncontroversially to possess. In order to explain the fact that legitimacy is taken for granted in a way that

obligating power is not, this uncontroversial property would need to be one from which obligating power does not also follow. The most plausible such line of thought appeals to the property of *justification* described in the introduction.[60] As we said before, a state's being justified in this sense is very roughly a matter of its *existence* being in some sense *a good thing* or *acceptable.*

As I also mentioned when I introduced the notion of justification, I think I am using the term 'justification' in roughly the same way as A. John Simmons, when he famously distinguishes justification from what he calls legitimacy (by which he means something different to me).[61] What Simmons explicitly says to characterise this notion is somewhat vague and it could be that he is in fact talking about something different, such as a state's being justified *in doing* what it does. But the phrases he uses generally seem to suggest that what he is taking to be justified when a state is justified is its *existence* (not its actions). For instance, he says that 'in the course of such a justification we will typically argue that certain virtues that states may possess or goods they may supply – such as justice or the rule of law – make it a good thing to have such states in the world'.[62] This seems quite clearly to suggest that what we are aiming to show is that such states' *existence* meets some moral bar, is in some sense *morally acceptable.*[63]

Of course, if we understood a state's 'justification' as its being justified in enforcing its directives (generally and exclusively), then justification would just become equivalent to legitimacy in my sense. To claim that justification in this sense is uncontroversial would just be to claim that legitimacy is uncontroversial, and this would not be much good as a defence of the safety of legitimacy. However, it might be thought that there is little or no distance between showing that a state's existence is justified (morally good or acceptable) and showing that the state is legitimate. Showing that a state is legitimate in my sense involves showing that it has a general and exclusive permission to enforce. It is essential, though, to something's being a state that it claim and exercise a general and exclusive permission to enforce. If a state ceased to make and exercise such a claim it would cease to be a state. So, does not showing that *a state* is justified in

---

[60] Simmons (2001) suggests sympathy for such a line of thought (131).

[61] Ibid. 122-37. Justification of states is discussed elsewhere and tends to be characterised vaguely. For instance, Schmidtz (1990) says that 'to justify an institution is, in general, to show that it is what it should be, or does what it should do' (90). Anticipating what is to come, it will be noticed that a characterisation like this just sidesteps questions about *whether* it should be or whether it should do anything.

[62] Ibid. 126

[63] Not everything he says points in this direction, such as when he talks about justifying a state by rebutting arguments that it 'practices wrongdoing' (Ibid. 124).

existing amount to showing that it has a general and exclusive permission to enforce? The answer to this is no. As I will explain in more detail below, it does *not* follow from some institution's existence being a good thing that it is morally *permitted* to do those things that are essential to it. It might be that the institution's existence is a good thing *despite* the fact that doing something impermissible is necessary to its existence.

Nevertheless, even if a state's justification (in my sense) is not simply *equivalent* to its legitimacy, it does have some plausibility as a candidate for a relatively uncontroversial property of some existing states (from which legitimacy might be thought to follow). It does seem at least *relatively* uncontroversial that there are some current states whose *existence* is a morally good thing. Why might we think this? There are a number of goods that states (or some states) achieve and that we should welcome. These include basic security, social coordination, distributive justice and more. Whether or not we think that *a* state or *this particular* state is *necessary* for the achievement of these goods, some actual states *do* achieve these goods, and it seems this is something we should be glad about. Further, it might be thought that for at least some states, the value of the goods that the state provides is greater than whatever disvalue it is responsible for (though this might be disputed). It might also be true, in many cases, that if we were to get rid of the existing state, things would be worse. So, it might be that for some states, on balance, their existence is something that is good, that we should be glad about. Whether any state's existence *is* justified will depend on what more precisely is meant by justification, or by a state's being something we should be glad about, and I will suggest below that there are several possible interpretations. But it seems plausible at least that some states will be morally good or acceptable in at least some sense. My question in this chapter will be what follows from this.

David Schmidtz distinguishes between teleological and emergent justification.[64] A teleological justification 'seeks to justify institutions in terms of what they accomplish', 'in terms of how they do or will serve [certain] goals', while emergent justification is 'an emergent property of the process by which institutions arise'.[65] Simmons's and my characterisation of justification is consistent with both of these. Schmidtz's two types of justification are different ways in which, or different senses in which, an institution might be morally acceptable. If a state is teleologically justified, it is morally acceptable in the sense that it is *good enough* in respect of some good it achieves. If it is emergently

---

[64] Schmidtz (1990)
[65] Ibid. 90-1

justified, it is morally acceptable in the sense that its coming to be did not violate any constraints of some kind.

Since a state is an exclusive enforcer, a state's coming to be involves its coming to be an exclusive enforcer. Unless a state is morally permitted to be an exclusive enforcer, its coming to be must violate some moral constraints. Thus, if a state is emergently justified, it will be legitimate. We can only be warranted in assuming that a state is emergently justified if we are warranted in supposing that it is permitted to be an exclusive enforcer, i.e., that it is legitimate. If a state lacks this permission then it cannot be emergently justified now, since by continuing to exist (as an exclusive enforcer) it will be violating some moral constraint. In what follows, then, I will consider only teleological justification.

It is common to talk about justifying *the* state (in general),[66] but, whatever exactly it means for *the* state to be justified, it does not obviously make any sense to talk about *the* state (in general) being legitimate. Thus, unless it is possible to argue from *the* state being justified to some particular state being legitimate, we will need to start from the claim that a *particular* state is justified. Simmons says that 'if "justifying the state" is to identify any plausible enterprise in political philosophy, then it should at least be taken to be accomplished if we can show that one or more specific kinds of state are morally defensible (comparatively or noncomparatively)'.[67] If all it takes to justify *the* state is to show that *some* specific kind of state is morally defensible, then it will be sufficient that the best possible (ideal) state be morally defensible (even if no other kind of state would be morally defensible). Simmons says that a justification 'will provide some comfort to those who have chosen to live in a justified state: their choice wasn't a dumb choice … nor was it a choice to participate in an immoral arrangement'.[68] If one's *particular* state is justified, then it will make sense to feel this way. However, if we learn that *the* state is justified in Simmons's sense, it is hard to see how this could provide any comfort at all to anyone living in states. If all it takes for *the* state to be justified is for *some* kind of state to be justified, then learning that *the* state is justified will not tell us state-dwellers anything useful except that states are not, by virtue of being states, *necessarily* morally indefensible. What *we* want to know is whether *our* state is justified. It is hard to see how the legitimacy of *this* particular state could follow from just *some* state's being morally defensible.

---

[66] As Simmons (2001) and Schmidtz (1990) do.
[67] Simmons (2001) 125
[68] Ibid. 126

Of course, it might be that the *kind* of state that we show to be morally defensible when justifying *the* state is a broader type than one only tokened by the best possible state, and this might be a type of which some actual states are tokens. However, if this is the case, we will want to know what it takes to justify a *kind* of state. If it just requires showing that *some* state of that kind is morally defensible, then knowing that a type of state of which ours is a token is justified will not necessarily tell us anything interesting about *our* state. If, on the other hand, it involves showing that *all* states of that kind are justified then it will be sufficient to show that *our* state (which is a token of the type) is morally defensible. However, the important thing is not that *the* state is justified, but that *this* state is justified. Of course, we can show that *this* state is justified by showing that all members of a set that includes this state are justified, but what is of interest to us is not the latter fact about the set of states, but the fact about this state that follows from it. Thus, I will focus on arguments that start from premises about particular states rather than *the* state.

What, more precisely, then, might the claim that a particular state is justified mean? What we have said already (in line with what Simmons says) is that justifying something involves showing it to be morally *acceptable.* But what, exactly, is it for something to be morally acceptable? Simmons provides some further characterisation of justification. He says that showing something to be justified 'centrally involves rebutting certain kinds of possible objections to it: either *comparative* objections – that other acts or institutions (etc.) are preferable to the one in question – or *noncomparative* objections – that the act in question is unacceptable or wrong or that the institution practices or sanctions wrongdoing or vice'.[69] He also says that 'we can justify *the* state by showing that some realisable type of state is on balance morally *permissible* (or ideal) and that it is rationally preferable to all feasible nonstate alternatives'.[70] Finally, he says that 'in the course of such a justification we will typically argue that certain virtues that states may possess or goods they may supply – such as justice or the rule of law –

---

[69] Ibid. 124. It might be thought that the distinction between justifications that meet comparative objections and those that meet noncomparative objections maps onto the distinction between teleological and emergent justifications, but I do not think this is the case. There may in principle be comparative objections to emergent justifications: it could be, for instance, that in order to be emergently justified an institution must have been the best available or not significantly worse than other available options when it was created. Similarly, there may be noncomparative objections to teleological justifications: it could be that if a certain institution does not secure some specific goal it is impossible for it to be justified.
[70] Ibid. 125-6

make it a good thing to have such states in the world'.[71] I have said already that emergent justification could not plausibly be the basis for thinking that legitimacy is on safe ground, so I will only consider ways we might understand 'moral acceptability' in teleological terms. Here are some possible ways we could understand the claim that state S is morally acceptable:

A. State S is not morally bad (or is good enough).
   Suppose that we can evaluate institutions according to their moral virtues and vices or the good or bad consequences of their existence. There could be an evaluative threshold below which an institution is just *bad*. Anything not bad is acceptable. (This threshold could be defined in terms of quantitative levels of goodness or in terms of certain specific constraints that must be met.)[72]
B. It is better (or at least as good) that state S exist than not.
C. It is permissible to create state S (or, it would be permissible to create state S if state S did not already exist). (Perhaps this goes along with the claim that it is permissible to sustain state S when it is in existence.)
D. It is impermissible to destroy or destabilise state S.[73]

Claims A to C seem like plausible readings of Simmons's moral acceptability idea and I cannot think of any other. Claim D seems less like a good interpretation of Simmons's notion of justification. However, since our purpose is not Simmons exegesis, but to identify a property states may possess that is both relatively uncontroversial and from which legitimacy follows, I include D to see if it might do this job. Claims C and D are of a different sort to claims A and B, but all concern the *existence* of state S, and could represent senses in which its existence is *acceptable*. Claims C and D, unlike A and B, concern agents' moral relation to the state's existence. I will argue that all of these interpretations of justification either fail to be uncontroversially true of some states, or can only license an argument to legitimacy with the aid of certain feasibility premises.

---

[71] Ibid. 126
[72] If the threshold to be met is understood deontologically, as the non-violation of specific constraints, it would be natural to understand a state's not being morally bad as its not doing anything morally impermissible. However, given that enforcement is something that all states do, and also, arguably, necessary for an institution to count as a state, to assume that a state is not morally bad in this sense would be to assume that its enforcement is morally permissible, i.e., to assume its legitimacy. Thus, *this* reading of 'is not bad' cannot be used in an argument for a state's legitimacy. The threshold of goodness will have to be understood either teleologically, or in terms of its meeting deontological constraints *excluding* the permissibility of its enforcement.
[73] It might seem that C and D here make justification an *emergent* property, but this would be a mistake. That it is permissible to create state S or would be permissible to create it if it did not already exist is not a feature of the history of state S. It has nothing to do with how state S came about or continues in existence. Rather, it is a non-backwards-looking property of a state and so could be a reading of what it is to be *teleologically* justified.

1. The argument

How, then, might we argue from these claims' being true of a certain state to that state's possessing a general and excusive moral permission to enforce? As mentioned above, we might be tempted to do so on the basis of a conceptual premise about states. As we have defined them, it is a conceptual necessity that states claim a general and exclusive moral right to enforce compliance with their directives, and that they actually have at least a somewhat extensive *power* of enforcement. Thus, it might be thought that an argument like the following can be made:

1. State S is justified (in one of the above senses).
2. It is an essential feature of state S that it claim and possess general and exclusive enforcement power.
3. Therefore, state S is justified in claiming and possessing general and exclusive enforcement power.
4. Therefore, state S is morally permitted to be an exclusive and generalised enforcer, i.e., it has a general and exclusive moral permission to enforce.

For interpretations A to C, however, this argument will not be valid. It does not follow from the fact that some entity is morally good, or that its existence is morally good (or better than its non-existence), that some necessary feature of it is also morally good. Consequently, it also does not follow that it is morally permitted to do something that is necessary to its existence. The entity in question might be morally good, or its existence might be a good thing, *in spite of* some necessary feature of it. A state might be morally good, or its existence might be better than its non-existence *despite* the fact that it necessarily does something impermissible. Similarly, it does not follow from its being permissible to create a state, that that state is morally permitted to do something that is a necessary feature of its existence. It might be permissible to create the state *despite* the fact that it will do something impermissible.

For interpretation D, though, it might be that some conceptual argument of this form will go through:

1. It is impermissible to destroy state S.
2. It is an essential feature of state S that it claim and possess general and exclusive enforcement power.
3. Thus, if state S ceased to claim and possess general and exclusive enforcement power, it would cease to exist.
4. Thus, state S's abandoning its general and exclusive enforcement power would amount to self-destruction.

5. Therefore, it is permissible for state S to exercise general and exclusive enforcement power, i.e., it has a general and exclusive moral permission to enforce.

This argument may be valid. I do not think, though, that it will do the job that we need of providing a bridge between some (relatively) uncontroversial property of some existing states and the legitimacy of these states. I will return to claim D and this argument below.

There is another form of argument, though, that might seem to license the inference from the other interpretations of justification to legitimacy. This form of argument could be thought to work for any of the suggested interpretations of justification. First, then, I will consider the general form such an argument might take leaving an interpretation of justification to be filled in, and then, below, I will take each interpretation in turn to ask whether such an argument can successfully establish legitimacy when justification is interpreted in each way. This more promising form of argument relies on the idea that a state's being an exclusive and generalised enforcer is required *for it to be justified*:

1. State S is justified.
2. State S's being justified *requires* it to be an exclusive and generalised enforcer.
3. Therefore, state S is justified as an exclusive and generalised enforcer.
4. Therefore, state S is morally permitted to be an exclusive and generalised enforcer, i.e., it has a general and exclusive moral permission to enforce.

Premise 2 should be read as claiming that the qualities of the state in virtue of which it is justified are only possible when the state enforces its commands or directives. The thought behind this argument could be put in the following way: state S is justified *because* of its exclusive and general enforcement and so its enforcement must itself be justified, and so permitted.[74]

Why might we think that premise 2 is true? That is, why might we think that a state must have general and exclusive enforcement power in order for its existence to be a good thing, or better than its non-existence, or for it to be permissible to create/impermissible to destroy? There are various roles that

---

[74] We said before that the strength of conclusion we are looking for, the strength that would count as a defeat of the anarchist, says that the state has a general and exclusive permission to enforce *at least* all of its just directives, whatever they may be. Thus, we can read the conclusion of the above as shorthand for this. Consequently, we will also need to read premises 2 and 3 as shorthand for the claims that state S's being justified requires it to have the exclusive and generalised power to enforce at least all of its just directives, and that this power is justified.

states fulfil, a number of which may seem to make their existence a good thing, or permissible to bring about. Some of these roles might be thought to require an exclusive enforcer in order to be successfully fulfilled. One of these is a coordinative role. States' establishment of artificial rules plays an important social coordinative role and makes possible all sorts of social cooperation that would not be possible otherwise. For instance, states establish traffic laws, they regulate exchange and so on. Another is a justice-ensuring role. There may need to be certain rules established to ensure that, for example, the distribution of benefits and burdens is just. This function could be performed by a state. In order for a set of directives to play either of these roles, they must in some sense be *dominant* in that society. If there are multiple competing institutions each issuing incompatible coordinative directives claiming to be authoritative and aiming to solve coordination problems, so long as there is not one whose directives are dominant (whose directives are followed in preference to the others), none will be able to fulfil the coordinative function. Similarly, if there are multiple competing institutions each issuing incompatible justice-ensuring directives (supposing there is more than one way in which justice can be achieved), it is likely that none will successfully achieve justice. Furthermore, since there are many different sets of rules that could play the coordinative and justice-ensuring roles, it might be thought that there needs to be a single *institution* (such as a state) that determines *what* the dominant rules will be. This institution will need to have the power to make the rules it chooses dominant, whatever they are. Thus, the thought will go, a state's being justified requires it to have the power to make its laws or directives dominant, at least so long as they are just (the power to make *unjust* rules dominant cannot contribute to a state's being *justified*).

Now it is possible (at least conceptually) that there could be one such institution that achieves this sort of dominance without the use of enforcement. There could be an institution that issues directives that are sufficiently well observed for them to play these functions but that makes no attempt to enforce them, and which either has no competitors or is generally recognised as dominant over whatever competitors it does have. How could an institution achieve this sort of dominance without the use of enforcement? This is a difficult sociological question. There is no *a priori* reason, though, why it should not be possible for this to happen. However, there is a popular sociological premise that supports the view that a state (or other institution) could not play these roles without enforcing its commands or directives. The sociological premise is the Hobbesian claim that human nature is such that it will not be possible to achieve universal adherence to

a set of rules among members of a society unless these rules are enforced.[75] It will not be rational for people to obey the directives established by some institution unless their doing so is *enforced.*

Further, it might well be thought that an *exclusive* enforcer is necessary to achieve universal adherence. If there are multiple enforcers, the situation is no different to how it would be if there were no enforcers. If it is not possible to achieve universal adherence to a single set of rules among members of a society without enforcement, then there seems no reason to suppose that it is possible to achieve universal adherence to (that is, enforcement of) a single set of rules among multiple enforcers. If there are multiple enforcers without any superior agent enforcing their compliance with a dominant set of rules, they are in just the same situation as the members of a society without any enforcement. Thus, we have an argument from justification to legitimacy that goes like this:

1. State S is justified.
2. State S's being justified *requires* it to determine and to make dominant a set of just rules.
3. It is only possible to establish a dominant set of rules that are generally adhered to in a society by exclusive enforcement.
4. Therefore, state S's being justified requires it to have the exclusive power to enforce all its just laws and directives, whatever they may be.
5. Therefore, state S is justified as an exclusive enforcer of all its just laws and directives.
6. Therefore, state S is morally permitted to be an exclusive enforcer of all its just laws and directives, i.e., it has an exclusive moral permission to enforce all its just laws and directives.

Premise 3 is stated as a strict possibility claim. However, it is clearly false on at least some readings of possibility. It is not logically or metaphysically impossible for there to be a dominant set of rules universally adhered to without there being any enforcement. I think that the most plausible reading of this premise is as a *feasibility* claim: that is, as the claim that it is not *feasible* to establish a dominant set of rules in a society without exclusive enforcement. Though it might be strictly *possible* (there is *some* possible world in which state S establishes a dominant set of rules without exclusive enforcement), the claim would go, state S

---

[75] The argument that comes below is a Hobbes*ian* argument in spirit, though it is certainly not *Hobbes's* argument. I think Hobbes has no argument for state legitimacy because he thinks none is needed; the sovereign's permission to enforce is just left over from the unlimited right to all things that all persons have when not subject to a sovereign power, as I argue in Guillery (forthcoming). However, the empirical premise seems to follow from Hobbes's arguments about human nature and the *reason* Hobbes takes the state to be justified is that it enforces its commands.

has no *feasible* alternative to being an exclusive enforcer (if it is going to establish a dominant set of rules). My aim in this first part of the thesis is to show that plausible arguments for state legitimacy depend on feasibility premises and here we have found that what seems to be the most plausible line of thought underlying the assumption that state legitimacy (for at least some actual states) is on safe ground depends on a feasibility premise.

However, this might be thought too easy. It might be said that, although there are *some* possible worlds in which a dominant set of rules is established without exclusive enforcement (it is not *logically* or *metaphysically* impossible), there is some restricted range of accessible possible worlds in *none* of which is a dominant set of rules established without exclusive enforcement. That is, it might be thought that there is a restricted sense of possibility in which it *is* impossible to establish a dominant set of rules without exclusive enforcement. Of course (similarly to what I will argue is the case with feasibility) there are a number of different sorts of possibility, a number of different ranges of possible worlds (accessibility relations) relative to which possibility claims might be made. It is not *obvious* which of these sorts of possibility are constraints on moral requirement (certainly they are not all). However, it might be claimed that certain possible worlds are *too* distant to be relevant to moral permissibility; that is, impossibility over a range of worlds that excludes these distant worlds straightforwardly rules out moral requirement. It could be thought that the only possible worlds in which a dominant set of rules is established without exclusive enforcement are this distant, and so that this is *impossible* for all intents and purposes (at least, for the purpose of determining moral permissibility). The claim, then, would be that, although it is not clear exactly which sorts of possibility (short of logical and metaphysical possibility) constrain moral requirement (and so whether premises like 3 above allow arguments about permissibility to go through), the possible worlds which make premise 3 false are *so* distant that they cannot be relevant to moral permissibility.[76]

I think that it is not obvious whether this is the case, but let us suppose for the sake of argument that it is. In this chapter, I hope to show that the only plausible

---

[76] In terms of the account of feasibility that I introduce in chapter 5, this thought could be construed as the thought that there are certain FCs that are so unrealistic that feasibility given *these* FCs is *obviously* irrelevant to moral permissibility, and the establishment of a set of dominant rules without exclusive enforcement is only feasible on such unrealistic (and irrelevant) FCs. (On my account, feasibility is not simply equivalent to possibility, but just as there are undoubtedly some possible worlds in which a dominant set of rules is established without enforcement, there are also undoubtedly sharpenings of feasibility on which it is feasible to do so.)

arguments from justification to legitimacy depend on other infeasibility premises that *cannot* plausibly be thought to be of this sort. These infeasibility premises are made false by possible worlds that are *not* so distant that they are *obviously* irrelevant to permissibility. To be clear, I think that the dependence of the above argument on premise 3 is already enough to show that its success should not be taken for granted without some argument about exactly what role feasibility (and possibility) can play in constraining moral permissibility or impermissibility. But I aim to show that *even if* the impossibility claim in premise 3 is thought to be so robust that it is *obviously* a constraint on moral requirement/permissibility, the argument from justification to legitimacy can only be successful with the aid of *other* infeasibility premises that are much less robust, and much less obviously a constraint on moral requirement/permissibility.

It could be thought that premise 3 should be read instead as a *probability* claim, that is, as the claim that only through exclusive enforcement is a dominant set of rules *likely* to be established. This will not work, though. Probability is not relevant to moral permissibility in the way that feasibility and possibility are ordinarily thought to be. The claim that 'φ-ing is the only feasible way of doing something morally required (or, perhaps, morally good, or what is the moral best)' seems, at least on the face of it, to be sufficient for the permissibility of φ-ing. However, we cannot substitute 'probable' for 'feasible'. That something is the only *probable* way for you to do something morally required or important does not show that you are permitted to do it. It could be that it is *you* that is making other possible ways of doing the required thing improbable. That you are unlikely to do those alternative things does not excuse you from doing them if there is otherwise a requirement to do those things rather than the thing that you are more likely to do. What we need to know is not whether it is *probable* that you will do the required thing in any other way, but whether there is any other *possible* or *feasible* way for you to do so. Perhaps it could be that something's being the only way of doing something required that would be likely to succeed *if you tried* is sufficient to show that it is permissible. On one account (one that I will argue is mistaken) this is what feasibility is: probability conditional on trying. My arguments in chapter 5 will show that this also is not a constraint on moral requirement and permissibility and so does not allow the kind of inference involved in the above argument.

The argument form I have just suggested leaves the precise meaning of 'justification' unspecified. I will now try out the four different interpretations of

justification suggested above (page 48) and see if any of them can get us a successful argument to legitimacy.

2. A

On this interpretation, to say that a state is justified is to say that it is not morally bad (or is good enough). Let us grant for the sake of argument that there is some plausible threshold of goodness for which this claim is sufficiently uncontroversially true of some possible or actual states. This interpretation of justification gives us the following argument:

1. State S is not morally bad (is good enough).
2. State S's not being morally bad counterfactually depends on its determining and making dominant a set of just rules.
3. It is only possible to establish a dominant set of artificial rules that are generally adhered to in a society by exclusive enforcement.
4. Therefore, state S's being not morally bad requires it to have the exclusive power to enforce all its just laws and directives, whatever they may be.
5. Therefore, state S's being an exclusive enforcer of all its just laws and directives is not morally bad.
6. Therefore, state S is morally permitted to be an exclusive enforcer of all its just laws and directives, i.e., it has an exclusive moral permission to enforce all its just laws and directives.

Now, if premises 2 and 3 hold, and if state S exists, then there are three possible ways state S might be:

X. State S is not bad and has and uses an exclusive power of enforcement (of its just laws and directives).

Y. State S is bad and has and uses an exclusive power of enforcement.

Z. State S is bad and lacks or does not use an exclusive power of enforcement.

Now, *if state S exists*, given 2 and 3, the only possible scenario in which state S lacks an exclusive power to enforce is one in which it is bad: it is not possible for it to lack this power and not be bad. Thus, if X, Y and Z were the only possibilities, it may seem reasonable to conclude that state S has an exclusive permission to enforce its just laws and directives. If the only way to avoid there being a morally bad state is for state S to φ, then plausibly it is permissible for state S to φ. (Even this is not obviously true, since if we are not consequentialists there might very well be certain things that we are not morally permitted to do even if they are necessary for bringing about certain desirable states of affairs, such as states of

affairs in which there is not a bad state; my point is just that we would need *at least* this to make a good argument of this sort for permissibility.)

However, X, Y and Z are not the only possibilities since state S *need not* exist. Thus, claims (5) and (6) do not follow from 4. The permissibility of enforcement for state S seems to follow if we *assume* that state S must exist, since then the only way to avoid there being a morally bad state is for it to be an exclusive enforcer. If we do not assume that, though, we allow the possibility that there are ways of not having a morally bad state that do not involve state S at all. In other words, it is not sufficient that state S's being an exclusive enforcer is necessary for *state S* not to be morally bad; we would need the claim that state S's exclusive enforcement is necessary for *things* (in general) not to be morally bad.

By way of (imperfect) analogy consider the case of a neighbourhood watch group. Suppose that a group of neighbours band together to form such a group with the aim of protecting their neighbourhood from crime and exacting punitive justice. Suppose that the exaction of justice is a *good* thing, in some sort of teleological sense: states of affairs in which justice is exacted are better than states of affairs in which it is not. Suppose also that the existence of the neighbourhood watch group is a good thing overall, in a similar teleological sense.[77] Whether or not its actions are permissible, it makes such a contribution to achieving security and justice (and perhaps performs other unrelated positive functions) that its existence is all-things-considered good. Now suppose that when the neighbourhood watch group is formed, its members come together to hunt down and exact justice on perpetrators of crime in their neighbourhood, but when they are together they egg each other on and become hungry for blood, so much so that if they do not exact justice on a guilty party (someone who in fact committed a crime), they *will* exact punishment on an innocent person instead (which, suppose, would make the group's existence no longer a good thing overall). Their hunger for blood is so strong that the only way for them to avoid hurting an innocent person once they are together is to punish a guilty person.[78] It does not follow from this that it is *permissible* for the group to exact justice on the guilty (even though we have supposed that the exaction of justice is a *good* thing). The group need not exist, so

---

[77] The existence of the group cannot be assumed to be good in a sense that includes the permissibility of its exacting justice, since this is what is about to be at question.

[78] It might be thought that this is inconsistent with the assumption that the existence of the group is a *good* thing. But suppose that the group is in fact very effective at identifying perpetrators of crime and there is sufficient crime that the group very rarely needs to satisfy its bloodthirstiness by punishing the innocent. We have also supposed that the exaction of justice is a good thing, so even in this scenario, the existence of the group remains (teleologically) a good thing.

it is not the case that their exacting justice on the guilty is the *only* way to avoid their hurting the innocent; they could instead disband.

Even if the exaction of justice is itself good, it might nevertheless be *im*permissible for the neighbourhood watch group to carry it out because, for instance, it is better that some other agent exact justice instead (perhaps the state or the community as a whole) or because the group simply has no right to do it. It is not plausible, I think, that (in Simmons's terms) permissibility can always be established simply by meeting all *noncomparative* objections. Actions (including creation of institutions) can be impermissible for wholly *comparative* reasons. On a maximising moral theory that only permits *optimal* action, this is particularly clear.[79] But even on a non-maximising theory, it seems perfectly plausible that actions can be impermissible for comparative reasons. Actions may be impermissible even though there are no noncomparative objections, simply because of what they make impossible.[80] (In the case of institutions, it is perfectly plausible that sometimes it will be impermissible to create an institution that is subject to no noncomparative objections simply because there are alternative institutions whose creation would be so much better and whose creation is made impossible by the creation of the first institution.) Thus, I think, the above argument is invalid.[81] A state's being an exclusive enforcer may be impermissible even though there are no noncomparative objections, simply because of superior alternatives that it makes impossible.

As we saw, though, the argument might seem to work if there is no possibility of state S not existing. (The argument then goes through if it follows from the fact that X's φ-ing is the only way to avoid there being a morally bad state that it is permissible for X to φ.) It may be, then, that we could make the argument work by adding the premise that it is not feasible for state S to stop existing (or that there are no feasible morally good worlds in which state S stops existing). This is

---

[79] Unless, of course, the maximising theory is interpreted as saying merely that we *ought* to do what is optimal, but not that we are *required* to.

[80] I am not here assuming much about what it takes to establish permissibility. I make no positive claim about what is *necessary* to establish permissibility, only about what is not sufficient. There seem to be cases where permissibility is ruled out for comparative reasons, so just showing that there are no noncomparative objections to something is not sufficient to *establish* that it is permissible or permissible to create.

[81] Note that what I have said here does not depend on whether you have a consequentialist or deontological view of permissibility. If you have a deontological view, it still might be that φ's being necessary for avoiding there being a morally bad state is sufficient for its permissibility. But, as the above shows, this is not what the argument from A gets us. For a consequentialist, on the other hand, presumably the only permissible actions will be those that bring about the *best* consequences, but a state's existence simply not being *bad* does not show that the *best* consequence is one in which it exists.

certainly not the sort of feasibility premise that is *obviously* a constraint on moral permissibility. Possible worlds in which a particular state stops existing (even morally good worlds in which a particular state stops existing) are not that distant and, especially since we are considering what is feasible *for the state*, it is not at all obvious that it is infeasible for it to stop existing. Thus, to warrant taking the success of this argument for granted, we would at least need to devote some attention to the concept of feasibility and when exactly it is a constraint on moral argument.

3. B

Interpretation B is the claim that it is better (or at least as good) that state S exist than not. Here again I will grant for the sake of argument that there can be some states for which B is sufficiently uncontroversial. However, it would be a mistake to move from B to the claim that state S is legitimate. It might seem that a version of the above argument can be made as follows:

1. It is better (or at least as good) that state S exist than not.
2. (1) is true *because* state S determines and makes dominant a set of just rules.
3. It is only possible to establish a dominant set of rules that are generally adhered to in a society by exclusive enforcement.
4. Therefore, state S's existence being better than not depends on it having the exclusive power to enforce all its just laws and directives, whatever they may be.
5. Therefore, it is better that state S be an exclusive enforcer of all its just laws and directives than not.
6. Therefore, state S is morally permitted to be an exclusive enforcer of all its just laws and directives, i.e., it has an exclusive moral permission to enforce all its just laws and directives.

I am not at all sure that (5) here follows from (1) and (4), but it is the inference from (5) to (6) that I will question here. It might also be thought that there is a different route to legitimacy: if B is true, then C is true (it is morally permissible to create state S, or sustain it in existence), and this gets us closer to legitimacy since at least it is in the realm of permissibility. (We will see below that it is more plausibly possible to get from C to legitimacy.) However, I think that the move from B to C is illegitimate, for closely analogous reasons to the reasons that moving from (5) to (6) is illegitimate. I will here focus on the move from a state's existence being better than not to the permissibility of creation (B to C), but the same considerations apply to the move above from a state's exclusive enforcement being better than its not doing so to the permissibility of exclusive enforcement.

Claim B says that the actual world in which the state in question exists is better than the world would be if that state did not exist (or if the state in question is one that does not currently exist, that the world would be better than it actually is if the state in question existed). Either way the claim is a counterfactual one: 'if state S did not exist, the world would be less good than it actually is' (or 'if state S existed, the world would be better than it actually is'). On standard (Lewisian) analyses of counterfactuals, this means roughly that the actual world in which the state exists is better than the closest possible world in which it does not.[82] The move that would have to be made goes from the state's existence being better than its non-existence to its being permissible to bring it into existence or sustain it in existence (or from the world being better with the state's exclusive enforcement than without to its being permissible for the state to be an exclusive enforcer). However, this is an illicit move, since it does not follow from its being the case that state of affairs S is better than one other counterfactual possible world that it is (or was) permissible to bring about S or sustain S. This is to narrow the comparison class too much. If some state of affairs S is better than *any* other possible world then this *might* plausibly be sufficient to show that it is permissible to bring about S.[83] But that S is better than *some* other possible way things might be, even if that is the closest possible world in which S does not hold, does not show that it is permissible to bring it about.[84] Suppose, for instance, that in the actual world I slap you in anger. Suppose also that in the closest possible world in which I do not slap you I stab you. It is clearly better that I slap you than that I stab you. In some sense, then, it is better that I slap you than that I do not. This does not mean, though, that it is permissible for me to slap you.[85]

---

[82] Strictly, on Lewis's analysis, it means 'there is a world W, in which the state does not exist and which is worse than the actual world, that is closer to the actual world than any world X, in which the state does not exist and which is not worse than the actual world'. Lewis (1973)

[83] Even this need not be sufficient, since if we are not consequentialists we will not believe that it is generally true that it *must* be permissible to bring about the best possible state of affairs, but it does seem plausible that some state of affairs' being the best possible presents a good prima facie case for the permissibility of bringing it about, while a state of affairs S simply being better than *some* other possible world (even the closest one in which S does not obtain) does not.

[84] Nozick (1974) makes a similar point, arguing that we cannot justify 'the state' by comparing it simply to the anarchic situation that would exist if the actual state situation did not (4-6). He, however, is interested in justifying *the* state in general, showing that it is desirable to have a state and that the existence of one need not violate anybody's rights, rather than in showing that any particular states are *legitimate*.

[85] This might recall the debate between actualism and possibilism in moral philosophy. Jackson and Pargetter (1986) famously defend 'actualism', the view that an agent *ought* to perform an act if and only if what would happen if the agent performed it is better than what would happen if she did not. There are strong objections to actualism (see, for instance, Wedgwood (2009) and Ross (2012)), but whatever the plausibility of actualism

Perhaps things will seem better if we proceed in the other direction, claiming that a state that does not currently exist would be better than the actual world. Here we simply say that we would improve the world by creating the state in question (assuming the costs of changeover are not too high). However, though it may sound plausible that any act that improves the world morally is permissible, this cannot be true. If, for instance, by improving the world in one way you make it impossible to improve the world even more, then it could be impermissible to improve the world in the first way. (There may also be rights considerations: if the improved world is just better *overall*, but not necessarily better for everyone, the rights of some could be violated in bringing about the better world.)

Now, maybe we could replace B with the claim that the existence of a particular state *is* in fact better than *any* possible world in which that state does not exist. Note that this must be the claim that for whatever state we are trying to establish a permission to enforce, at least some possible world in which it does enforce is better than any possible world in which *that* state does not exist. It must be better than any other state. Otherwise we do not have the claim that *this* state's existence is better than any other way things might be. Thus, such a claim could only establish the permissibility of enforcement for the best possible state. Furthermore, the claim seems quite plausibly not to be true of any state. Is the best possible state world better than the best possible non-state world? This is debatable, but the best possible non-state world is presumably one in which people live together in society harmoniously without disputes and flourish communally and freely and in which there is no state with a generalised power of enforcement, while in the best possible state world there is a state with a generalised power of enforcement. There is nothing obviously deeply (logically or physically) *impossible* about the former world: at least some argument would be needed to show that there is. It seems to me that the former is obviously more desirable.

Note that my claim is not that for some action x to be permissible it is *necessary* for there to be a possible world in which x occurs that is better than all possible worlds in which x does not occur. This would be extremely demanding. Rather my claim is that the only straightforward way of *establishing* moral permissibility through comparisons of the goodness of alternate states of affairs is by making

---

as a principle about 'oughts', it is surely not plausible as a principle about *permissibility*. It is not the case that it is *permissible* to φ whenever what would happen if you did not φ is worse than what would happen if you did, as illustrated by the above case.

this radical claim. Just to point out that a world in which x occurs is better than *some* other possible world(s) is not sufficient.

The best solution to this problem is perhaps to claim that the existence of a particular state is better than any possible world *out of some restricted range*. We could replace B with the claim that the existence of some state is better than any *feasible* world in which that state does not exist. This seems like it might make the argument work. (Certainly, it would if there is a straightforward "'ought' implies 'feasibility'" constraint and if it is always permissible to do the best thing available.) Again, though, this argument now depends on a feasibility premise that is not *obviously* of a sort that constrains moral argument. Possible worlds in which there are better alternatives to *this* state's exclusive enforcement are not *that* distant that they are *obviously* irrelevant to moral argument. Thus, again, some thought about feasibility and its role in moral argument is needed to warrant taking the success of this argument for granted.

Another version of claim B states that the actual world in which state S exists is better than things would be if state S were *torn down*. Then we do not have to compare the actual world to all other possible worlds but only to worlds in which state S is torn down, all of which would involve a significant amount of upheaval. We could then add the premise that exclusive enforcement is necessary to prevent the state from being torn down to get the conclusion that it is permissible for state S to be an exclusive enforcer of its directives (because if it was not, it would be torn down, and this would be worse). This argument could obviously only get a permission for existing states, but this is all that is really needed. However, again the natural reading of the claim is as a counterfactual, making the comparison only to the closest possible world in which the state is torn down. Again, this is not enough to establish any sort of permissibility. The claim then could be that the actual world is better than *all* of the possible worlds in which state S is torn down, but this seems highly implausible. There are some possible worlds in which state S is torn down peacefully and replaced with something much better. More plausibly, again, the claim might be that the actual world is better than any of the *feasible* scenarios in which state S is torn down. But again, whether this is true will depend on what exactly is meant by 'feasible' and so some thought will be needed about whether the kind of understanding of 'feasible' that makes this claim true also allows an argument of this form to go through.

4. C and D

i

Interpretation C understands the claim that a state is justified as the claim that it is permissible to create it, while on interpretation D, it is the claim that it is impermissible to destroy or destabilise it. These claims seem to stand a better chance of allowing an inference to legitimacy, being already in the domain of permissibility. For claim C, we get an argument that goes like this:

1. It is permissible to create state S (or, it would be permissible to create state S if it did not exist).
2. It is only permissible to create state S if (and because) it establishes (would establish) and makes (would make) dominant a set of just rules.
3. It is only possible to establish a dominant set of rules that are generally adhered to in a society by exclusive enforcement.
4. Therefore, it is permissible for state S to be an exclusive enforcer of its just laws and directives, i.e., it has an exclusive moral permission to enforce all its just laws and directives.

This argument goes through if we accept the plausible-sounding principle:

P. If it is permissible to create an institution N *because* it does φ, then it is permissible for N to do φ.

Let us accept, then, that we can get from C to legitimacy.

There is no plausible argument of the above form starting from claim D (the claim that it is impermissible to destroy or destabilise state S), on the other hand, since it is not plausible that it must be permissible to do what makes the destruction of an institution permissible: if destroying an institution is *permissible*, then why should we be concerned about it morally? We saw above, though, that there may be a *conceptual* argument from D available. (A state's ceasing to be an exclusive enforcer would be its ceasing to be a state, and so if it is impermissible to destroy it, perhaps it is impermissible for it to cease to be an exclusive enforcer.)

There is also a possible line of argument from D that starts from the following premises:

1. It is impermissible to destroy or destabilise state S.
2. A state is only stable if it establishes and makes dominant a set of just rules.
3. It is only possible to establish a dominant set of rules that are generally adhered to in a society by exclusive enforcement.

It may follow from these premises that it is impermissible to prevent state S from enforcing its just laws or directives. If obstructing a state's generalised enforcement will make the state unstable, then it is presumably impermissible by (1). The permissibility of enforcement for state S does not straightforwardly follow. There are all sorts of reasons why it might be impermissible to prevent a state from enforcing its directives. Some of them are perfectly consistent with the state in question lacking a general permission to enforce compliance with its directives. It could be that by enforcing its directives state S is violating a moral requirement, but that nevertheless, because preventing it from doing so would do more harm than good, prevention is impermissible. Similarly, it could be that state S lacks the virtues necessary for it to be permissible to create such a state when it does not exist, but now that it exists, preventing it from enforcing would do more harm than good, and so is impermissible.

One might argue, though, that (1) should be interpreted as applying to the state itself. If (1) holds, the thought goes, it is impermissible for the state to destabilise itself. If being an exclusive enforcer is necessary for stability, then it must be impermissible not to be. If it is impermissible not to be an exclusive enforcer (obligatory to be one), then it must be permissible to be one. There seems to be something odd about this argument. For one thing, it has the unexpected consequence that not only is it permissible for the state to be an exclusive enforcer, it is *obligatory*. Nevertheless, I am not sure what exactly, if anything, goes wrong with the argument, so let us suppose, if only for the sake of argument, that it is successful. Additionally, as we saw above, there is a *conceptual* form of argument from D to legitimacy that might be successful (which was not successful for the other interpretations of justification). There is at least a plausible case to be made, then, that legitimacy *does* follow from justification when interpreted as either C or D.

ii

However, these interpretations of justification fail to fulfil the other role required to justify the assumption that legitimacy is on safe ground. Interpretations A and B were plausible candidates for claims that are somewhat uncontroversially true of at least some existing states, but fail to warrant an inference to legitimacy. Interpretations C and D, on the other hand, may successfully get us arguments to legitimacy, but are not plausible candidates for claims that are uncontroversially true of some existing states. What we need is a claim that is more uncontroversial

for some state than the claims that the state is legitimate or has obligating power. Claim C, the claim that it is permissible to create a state, seems to be just as problematic or difficult to establish as the claim that a state has a generalised and exclusive permission to enforce. Creating an institution with a generalised and exclusive capacity to enforce raises the same (or similar) moral problems as does the possession or use of this generalised capacity.

It could be that C follows from some other uncontroversial claim. However, I have argued above that it does not follow from B. For similar reasons, C does not follow from claim A either. The reason C does not follow from B was that state S's existence simply being better than some particular alternative scenario is not sufficient to *establish* that it is permissible to create state S, because there might be some other possible scenarios that are so much better that it is in fact impermissible to bring about S's existence. Similarly, that state S is not morally bad (meets some evaluative threshold) does not rule out its being impermissible to create state S, because there could be alternative better scenarios that creating state S would prevent from coming about. That φ-ing prevents something better from happening is not in itself sufficient to show that φ-ing is impermissible, but it seems plausible that there *could* be cases where the fact that φ-ing would prevent something so much better makes it impermissible. To assume that state S's not being morally bad *shows* that it is permissible to create it does not allow for this. Thus, if it is still thought that C follows from another more uncontroversial claim, some plausible other candidate would need to be found.

One way of interpreting contractarian arguments for the state is as offering just such an argument for C. Jean Hampton, for instance, argues that a social contract argument (of the sort espoused by, notably, Hobbes, Locke and Kant) is supposed, amongst other things, '1. to explain the state as an entity whose origination and continued existence are the responsibilities of human beings, 2. to show why human beings are justified in creating and maintaining a state ...'.[86] The use of 'justified' here, if interpreted morally, seems to mean nothing other than 'permitted'. The contractarian argument is supposed to show that we *need* a state, and thus that we have some compelling reason to meet that need, i.e., to create or maintain a state. This is the point, for instance, of Hobbes's state of nature argument: human nature is such that without a state, human life will be a state of war, so for the sake of peace (which we are presumed to want) we need a state. For Hobbes the compelling reason to create or maintain a state is self-interested,

---

[86] Hampton (1986) 269

for Locke and Kant it is moral. The moral version of this argument will have as its conclusion the claim that it is morally permissible to create or maintain a state.

Interpreted thus, the contractarian argument is essentially an attempt to make the argument from B to C that I have argued against above. Instead of simply comparing the world with some given state to the closest possible world without *that* state, this argument compares it to (presumably) the closest possible world without *any* state. The problem is that we need to establish legitimacy for some specific actual state, so to establish permissibility of creation we would need to show that *this* state is better than all the relevant alternatives to *it*, not only the alternatives to there being *any* state. Even if a contractarian argument could show that the closest possible non-state world is worse than *all* state worlds, this is not sufficient to show that it is permissible to create any *particular* state. Suppose, for instance, S and T are both states that could be equally feasibly created from a state of nature, and though better than the state of nature, an S-world would be much worse than a T-world. It might well be in such a scenario that it is *not* permissible to create state S if by doing so we make it impossible to create state T. This is just the same problem as raised above. Thus, I think, even if legitimacy does follow from C, this is not sufficient to warrant the assumption that the legitimacy of actual states is on safe ground, for the simple reason that we are not warranted in assuming that the truth of C *itself* is on safe ground.

The case is similar with claim D, the claim that it is impermissible to destroy or destabilise state S. Again, we should ask why it might be thought that this is true of some actual state. The only obvious candidate response is some form of the thought that it would be *worse* if the state in question was destroyed or destabilised. This amounts to the same as the claim above that the actual world in which state S exists is better than things would be if state S were torn down.[87] That the actual world is better than the closest possible world in which the existing state is torn down does not seem to be enough to show that it is impermissible to tear it down, since there may be another close, but not quite as close, yet perfectly accessible world in which the state is torn down that is *better* than the actual world. On the other hand, it is very unlikely that the actual world with the existing state is better than *all* possible worlds in which that state is torn down. Thus, the impermissibility of destruction claim stops looking like such an

---

[87] Claim A won't help either, even more clearly than in the previous case. That something is not morally bad, or meets some evaluative threshold, is not sufficient to show that it is impermissible to destroy it. Often, it is perfectly permissible to destroy something good in order to replace it with something better.

uncontroversial starting point. What might support it (as we saw above with permissibility of creation) is the claim that the actual state of affairs in which the state exists is better than all *feasible* worlds in which it is torn down. This latter claim, though, yet again requires some substantive argumentation and consideration of what sort of notion of feasibility is relevant to permissibility.

B. Explicit arguments for state legitimacy

The line of thought that I have just been discussing, one that takes legitimacy to follow from justification (which might in turn be thought to hold for at least some existing states), is one that I think often explains the commonly made implicit assumption that the legitimacy of at least some existing states is on safe ground. It is a line of reasoning that might seem to warrant *not worrying* about the existence of legitimate states. If you think that an existing state is *justified*, some form of the above line of thought might seem to justify not worrying. Though such an argument is not always made explicitly, I think that it is often an implicit thought of this nature that leads to the assumption that legitimacy is safe.

The legitimacy of some existing states is not always taken for granted, however. Many writers do attempt to give arguments to support the legitimacy of certain kinds of states (often supposed to include some actual ones). There is a very long tradition in political philosophy of attempting to find grounds for the justification of the state, political obligation and state legitimacy. These questions are not always distinguished, and even if the distinction is not missed, it is sometimes thought that the answers to them will be the same.[88] Nevertheless, there are a number of arguments that have been made that do offer grounds for the state's general and exclusive permission to enforce, whether bundled together with obligating power and justification or given bespoke treatment. I will now very briefly take a look at what I take to be the most prominent families of arguments for state legitimacy and I will argue that, just as with the sort of reasoning discussed above, they can only successfully establish legitimacy for existing states with the aid of feasibility premises. These arguments do not tend to make any claim to be independent of feasibility considerations, so this will not be to show that they are mistaken. However, the account of feasibility that I will give in chapter 5 has the result that arguments that rely on feasibility premises in this way cannot be taken for granted. The first family of arguments, contractualist and

---

[88] Some arguments in the literature that describe themselves as being about 'legitimacy' may not address legitimacy *in my sense* (for instance, where 'legitimacy' is used in a sense where obligating power is a necessary or sufficient condition). (This is not a *fault* of theirs; I do not claim that my use of 'legitimacy' is in any sense the 'proper' or 'best' one.)

necessity-based arguments, are of broadly the same form as the line of thought just discussed. They attempt to argue from the state's existence being in some sense *a good thing* to the permission to enforce. They depend on feasibility premises for the same kind of reasons as above. After discussing this first family, I briefly turn to arguments that focus on democracy, which, I argue, either fall into the first family, or are not really arguments that address the legitimacy sceptic (they defend a *necessary*, but not a sufficient, condition for legitimacy). The second family of arguments, based on consent, is no good to establish the legitimacy of actually existing states, I will argue.

1. Contractualist and necessity-based arguments

There is a family of arguments for state legitimacy that, like the line of thought considered above, attempt to argue from justification to legitimacy. These arguments are all of roughly the same form as those discussed above. They all argue that the state provides something *good*, or that its existence is in some way good or necessary, and that the state's exclusive enforcement is necessary for the state's providing the good, or being good, and conclude that the state has a general and exclusive *permission* to enforce. These arguments offer various ways to bridge the gap between the state's exclusive enforcement being necessary for something good, to the state's being *permitted* to be an exclusive enforcer. The arguments I looked at above argued from some particular state's enforcement being necessary for that state's being good. I argued that we cannot get straight from this to the permissibility of that state's being an exclusive enforcer, because there might be alternatives to *that state* that are also good. I did not question, however, the validity of a possible inference from a state's exclusive enforcement being necessary for *the world's* being good to its legitimacy. I suggested that, since such a strict necessity claim is unlikely to be true of any particular existing state, the argument requires a feasibility premise (this state's exclusive enforcement is the only *feasible* way for the world to be good) to go through.

Even the inference from something (this particular state's exclusive enforcement) being necessary for the world to be good, or the best way it might be, to its permissibility, however, is questionable. There are various different arguments that provide possible ways to bridge *this* gap. However, since my arguments in this chapter did not rely on questioning *this* inference, finding a way to bridge this gap will not change what I said above. For the same reasons as the arguments examined above, contractualist and necessity-based arguments that attempt to bridge the gap between necessity and permissibility will only be successful with

the aid of certain feasibility premises. This is not as such an *objection* to these arguments, since they do not generally claim to be a priori, or independent of feasibility considerations, and indeed, some make their reliance on feasibility premises quite explicit.[89] However, it does raise the question, as with the arguments above, whether we can understand the notion of feasibility in a way that licences the necessary inferences.

Arguments of this form are quite various and account for a good proportion of the explicit arguments for state legitimacy that have been made. Contractualist arguments which defend state legitimacy on the grounds that a principle permitting the state's exclusive enforcement could not reasonably be rejected, or would be agreed to in certain idealised conditions, are of this sort.[90] The *reason* that such a principle is reached through contractualist reasoning is that the state's exclusive enforcement is good, or necessary for some good, and the contractualist framework provides a way to bridge the gap between this and permissibility. David Copp's Humean argument offers another way to do the same.[91] He argues that those standards that 'actually function as well as can be to make things go well in society' are justified, and state exclusive enforcement makes things go better.[92] Rolf Sartorius's Anscombe-inspired argument claims that positive rights (based in customary morality) become moral rights 'when they involve what is necessary for the successful carrying out of a task that must be performed for the benefit of those on whose behalf authority is exercised'.[93] Christopher Wellman argues, on the basis of a Samaritan duty to aid others, that since the state's exclusive enforcement is necessary for peace and security, we have a duty to support (or establish) state enforcement, and so we have no right against state

---

[89] See, for instance, Wellman (1996) 217.

[90] Ladenson (1980) suggests a contractualist argument for state legitimacy (140). Rawls in *Political Liberalism* (1993) might appear to be making an argument of this sort: state enforcement is acceptable if citizens can reasonably be expected to endorse it (see pp. 136-7). If this is what he is saying, there is no reason to think that citizens of a state *can* reasonably be expected to endorse its claim to exclusive enforcement rights unless there is no feasible better alternative. However, I think Rawls is better read as taking the existence of the state as a starting point, and asking what is the most acceptable way for its exclusive enforcement power to be exercised. His answer to this question is that it must be exercised according to principles that reasonable citizens could reasonably be expected to endorse (136-7, 217). If we take state enforcement for granted in this way, though, we will not be able to show that state enforcement *is* justifiable, again, unless we think there is no, or no good, feasible alternative.

[91] Copp (1999). Hume's (1994, Essay 23) own argument is focused on political obligation and 'the virtue of allegiance', but can easily be extended to legitimacy.

[92] Copp (1999) 37

[93] Sartorius (1999) 148. Anscombe's (1978) argument seems to be roughly the same; Sartorius's differences with her seem to concern political obligation, not legitimacy, but her paper is difficult to interpret and so I am not confident that this is exactly her argument.

enforcement.[94] Allen Buchanan argues, similarly, that a state's exclusive enforcement does not violate any rights since we all are under a 'robust natural duty of justice', the fulfilment of which requires the existence of an exclusive enforcer.[95] This latter claim seems best understood as a feasibility claim. Finally, Massimo Renzo offers a similar argument from the necessity of the state for order and security, only basing legitimacy in a right of self-defence, rather than positive Samaritan duties or duties of justice.[96]

As already noted, all of these arguments are of roughly the same form as the arguments discussed above. Consequently, they too all depend on feasibility premises. Their claims that state enforcement is *necessary* for some good, or for the world to be good, or for some moral requirement to be met will depend on the premise that there is no feasible way of achieving the good or making the world good or meeting the moral requirement without state enforcement. Additionally, to successfully establish the legitimacy of any actual existing states, these arguments will need to not only show that *the* state is necessary for fulfilling some task, but also that some *particular* state is too. In order to do this, they will need something like the premise that *this particular state's* exclusive enforcement is the only feasible way for the good in question to be provided (for the task to be fulfilled, in Anscombe's terms). Feasibility premises seem needed both to show that *the* state (some state) is necessary and to show that a *particular* state is necessary. (The feasibility premises required for the latter are less plausibly strict constraints on moral argument.)

2. Democracy

There is a large literature on the justification of democracy. This is sometimes, though not always, presented as involving, at least in part, showing that democratic states (states that are sufficiently democratic in the right respects) are *legitimate* in my sense.[97] It is worth noting, though, that justifying democracy *need not* involve this at all. There are at least two other sorts of questions, answers to which might be thought to amount to justifications of democracy. First, we might simply ask why, or whether, we should have democratic institutions or why we should use democratic procedures for making collective decisions. Second, we

---

[94] Wellman (1996) and (2009)
[95] Buchanan (2002) 707-8
[96] Renzo (2011)
[97] Estlund (2009) and Kolodny (2014a) and (2014b), for instance, both present their defences of democracy as *including* a defence of the legitimacy of democratic states (though both are also supposed to establish the *authority* of democratic decisions). (They both use the term 'legitimacy' in roughly the same way as I do.)

might ask why, or whether, a decision's being made democratically gives anyone a *reason* to implement or comply with it. Answering neither of these questions amounts to justifying a democratic *state* (supposing we understand 'state' as I proposed to in the introduction, where a necessary condition for statehood is its claiming, and exercising to a reasonable extent, an exclusive right to enforce over some population). Democratic institutions need not be states. Democratic decision procedures could in principle be followed without any enforcement whatsoever. We might well have reason to create (or maintain) democratic institutions or to implement or comply with democratic decisions independent of anyone's having any reason to *enforce* compliance with democratic decisions. Thus, a justification of democracy might answer one of these other questions, without giving any argument for the *legitimacy* of democratic *states*.

Many, though, who seek to justify democracy do wish to show that democratic states are legitimate.[98] We can crudely divide justifications of democracy into two kinds: *outcome-oriented* and *procedural.* Not all accounts will fit neatly into one of these categories. I do not have space to go into the details of all the available justifications of democracy, but I think some sort of hybrid of the comments I will make about each variety of justification will apply to other forms. Outcome-oriented accounts attempt to justify democracy (answer any or all of the three questions distinguished above) by appeal to some feature of the *outputs* of democratic decision procedures. They might argue, for instance, that democracy is justified because it tends to produce the morally best, or just, decisions, or because it produces decisions that correspond in the right way to the preferences of those subject to the decisions, or because it produces decisions that people are likely to accept, and so on.[99] Procedural accounts attempt to justify democracy by appeal to intrinsic features of democratic decision procedures themselves. They might argue, for instance, that democratic decision procedures are essential to social equality, or to achieving equal opportunity for political activity, or that non-democratic procedures (or the failure to follow democratic procedures) express a negative judgment about some, and so on.[100]

I think that outcome-oriented justifications of democracy, when considered as arguments for the legitimacy of democratic states, have the same kind of form as

---

[98] For a view that, on the contrary, separates democratic provenance and legitimacy (in my sense), see Stemplowska and Swift (2018).
[99] There are some clear examples of outcome-oriented arguments in, for example, Mill (2015) 180-405 and Arneson (2002).
[100] Two examples of clearly procedural arguments: Singer (1973) and Kolodny (2014a) and (2014b).

the arguments discussed in the previous section. They claim that using democratic decision procedures is the only, or the best, way of producing some morally desirable outcome. If the relevant outcome is sufficiently important, this will be enough to show that we ought to want democratic institutions, or that we ought to use democratic decision procedures. It does not yet give us any reason to want a democratic *state*, since democratic decision procedures could be used without a state. If we add the claim that an exclusive enforcer, of a democratic kind, is necessary for the desirable outcome in question (or the best way of producing it), we have the same kind of argument as discussed above. This will only support the claim that a *particular* democratic state is legitimate if it is claimed that only through this *particular* democratic state's exclusive enforcement can the desirable outcome in question be produced (or that it is the best feasible way of producing it).

Procedural justifications of democracy, on the other hand, do not provide an answer to a legitimacy anarchist. They should rather, I think, be understood as defending a *necessary* condition for state legitimacy. (That is, they should be understood this way if they are understood as addressing the legitimacy question, and not just the institutional question or the authority question). Procedural arguments claim that democratic decision procedures are necessary for something of moral value, or to avoid violation of certain moral constraints. In Kolodny's account, for example, democracy is held to be necessary for legitimacy (as well as desirable and authoritative) because it is necessary for avoiding relations of social superiority and inferiority in society. A *state*, however, is not necessary for democratic procedures to be followed, and so this does not in itself give us a reason to desire a democratic *state*, or to think that such a state is legitimate. A procedural argument may give us reason to think that *only* democratic states can be legitimate, but it will not in itself support the view that any states *are* legitimate. Again, it could be *thought* to establish legitimacy because something like the above argument is assumed. That is, it might be assumed that the existence of a *state* is the only *feasible* way of achieving the procedural value of democracy or some other morally important good. But to show that any *particular* state is legitimate, this would need to be the claim that *that state's* exclusive enforcement is the only (or the best) feasible way of achieving the relevant moral value or requirement.[101]

---

[101] One of the most sophisticated and influential justifications of democracy, Estlund (2009), combines outcome-oriented and procedural elements, but it similarly offers an argument for the legitimacy of particular democratic *states* only on the assumption that

3. Consent arguments

The other main family of explicit arguments for state legitimacy are based on consent. If the subjects of a state *consent* to its exclusive enforcement power over them, it may be plausible to think, then whatever moral presumption there is against state legitimacy is defeated.[102] However, I think that none of the available consent arguments can show that actual existing states are legitimate, that is, none can defeat the (a posteriori) legitimacy anarchist.

The classical consent argument is the Lockean one.[103] This is an actual consent argument: it is argued that certain existing states are legitimate thanks to the actual consent they receive from their subjects. This view makes legitimacy individualised: states are legitimate with respect to those subjects who do consent and not with respect to those who do not. Since no state receives the *explicit* consent of many of its subjects, Locke demanded instead *tacit* consent. Tacit consent, he thinks, can be given by, for instance, using the roads a state provides, or simply remaining within the state's territory.[104] However, the problems with such an argument are familiar. As Hume famously pointed out, subjects of a state rarely have any realistic alternative to remaining in their state, and an action cannot plausibly count as valid consent if there is no realistic alternative to it.[105] I will not rehearse the details of this debate, but it is, I think, widely accepted that there are no actually existing states whose legitimacy can straightforwardly be grounded in universal actual consent: the kind of actions that count as valid consent are not performed by many, let alone all, of the citizens of any existing states.[106]

Robert Nozick offers a more sophisticated form of consent argument.[107] He argues that a state could emerge from a state of nature without any violations of rights (and would be likely to emerge if inhabitants of the state of nature are acting rationally). It is somewhat obscure what exactly the methodological import of this argument is supposed to be. On the face of it, the fact that a state *could* emerge

---

state enforcement is *necessary* for some reason. Estlund's qualified judges *could* reasonably reject the *exclusive enforcement* of a democratic decision procedure if it were possible (and feasible) to establish and make effective a democratic decision procedure with the same epistemic and procedural qualities without state enforcement.

[102] This is not uncontroversial, but we can suppose for the sake of argument that this much is true.

[103] Locke (1980) ch. 8.

[104] Ibid. s. 119

[105] Hume (1994) 193

[106] Simmons (1991) defends a version of this basic Lockean consent argument, but is comfortable with the anarchistic consequences.

[107] Nozick (1974)

without rights violations offers no argument at all for the legitimacy of any actual state (nor for political obligation in any actual state). However, Jonathan Wolff suggests what seems like the strongest interpretation of this argument, which shows how Nozick's hypothetical story can be relevant to actual states that did not emerge in this way.[108] According to him, Nozick's argument proceeds in two steps. The first step involves showing that rational people have good reason to consider themselves bound to the state, and thus, that they should consent to its rule and enforcement. This is done by showing that the existence of the state is a good thing (in other words, it is justified). Nozick's approach to doing this is to show that something like a state (a dominant protection agency) would emerge as a result of people's rational free action in the state of nature. Of course, considerations like those above suggest that the fact that the existence of *a* state is better than the state of nature is not sufficient to show that one ought to consent to one's *particular* state.

But it is the second step of Nozick's argument that is most important for our purposes. This step aims to 'show that even those who fail to be convinced by the argument for the state are, nevertheless, still morally bound to obey it, and therefore forcing their obedience violates no rights'.[109] It is this step that, if successful, will show that states (at least, those of a certain kind) have a general and exclusive permission to enforce *even where their subjects have not consented* to their doing so. The argument for this step is intricate, and I do not have space to go fully into its detail. But to sum up roughly what is important for our purposes, the argument is that when a dominant protection agency emerges from the state of nature, thanks to a large number of people buying its services and consenting to its enforcement (Nozick argues that the market for protective services in a state of nature is something like a natural monopoly), it comes to be permissible for it to determine and enforce what procedures for rights protection even non-members may use. The reason for this is that it is assumed that everybody has a right to protect themselves against infringements of their own rights, and it becomes permissible for the dominant protection agency to enforce its members' rights when they consent to its doing so on their behalf. Thus, although non-members of the dominant agency have a right to enforce their own rights, members have a right not to be subjected to an unfair procedure. The dominant protection agency thus has a right to protect its members from being subjected to an unfair procedure, and so it can permissibly enforce non-members' using only

---

[108] Wolff (1991) 47-52
[109] Ibid. 51

procedures that *it* deems to be fair. Thus, in effect it ends up with an exclusive and general permission to enforce (or to determine who may enforce what).

However, even if we grant that this shows how a state that has the consent of the majority of its subjects can be permitted to enforce its directives even on those who do not consent (and grant also that we *ought* rationally to consent to the state), it still does not show that any *actual* states are legitimate. Even if we all *ought* to consent to our states, it remains the case that very few people actually have. No existing states are in fact in the position of a dominant protective agency. Unlike the hypothetical dominant protective agency in Nozick's story, actual states have no claim to the self-defence rights of the majority of their subjects, because they do not in fact have their consent, and so cannot ground a permission to enforce the compliance of non-consenters in a permission to enforce the self-defence rights of the majority on their behalf. That the majority *would* consent if rational (or even if asked) is neither here nor there.

Some interesting, more sophisticated consent-based (or similar) accounts of legitimacy have been given more recently in a 'political realist' vein, notably by John Horton and Amanda Greene.[110] Both of these views are presented as *accounts of legitimacy* and I want to take no stand on whether they are plausible as such. It may well be that something like the sophisticated consent-based (or similar) accounts presented by Horton and Greene is the correct account of the concept of *legitimacy*. My use of the term 'legitimacy' was not meant to suggest an attempt to give the correct account of the ordinary concept, but was merely stipulative. I think that neither of these arguments, though, helps respond to the legitimacy anarchist *in my sense*, that is, the denier of existing states' possession of a general and exclusive moral permission to enforce.

Horton does seem to present his argument as including an argument for the permission to enforce. He holds roughly that states are legitimate if and only if they meet the criteria that their subjects hold to be relevant for legitimacy.[111] Although he denies that his account is a consent account (since what makes states

---

[110] Horton (2012) and Greene (2016). Williams's (2005) 'basic legitimation demand' is quite similar to these accounts, although it is unclear to me whether meeting it is supposed to be sufficient for legitimacy, and whether legitimacy is understood in anything like my sense (for this reason I focus on Horton and Greene). As Estlund (2017c) points out, once we distinguish it from 'the obvious fact that obstacles are obstacles and cannot be ignored' (the idea that it is *pragmatically* necessary for a political order to make sense to its subjects), it is somewhat unclear what reason there is to suppose that the legitimacy, in my sense, of a political order 'is nothing but its de facto acceptability to whatever freely formed points of view are extant, however morally bad they might be' (375). (This objection will apply also to Horton, if not to Greene.)

[111] Horton (2012) 141-2

legitimate is not the fact that their subjects consent, but their meeting the criteria the subjects consider relevant), on the most natural reading, like consent theories, his account seems to have the consequence that no actual states will be legitimate (at least not generally so). This is because, almost inevitably, in actual states the criteria that people hold to be relevant to legitimacy vary, and existing states are not likely to meet them all. States may, then, be legitimate with respect to some of their subjects, but their *universal* claims to legitimacy with respect to all their subjects will not be vindicated.

Unlike Horton, Greene does not obviously present her account as offering the grounds for permissible state enforcement. Thus, it need not be to disagree with her to state that it does not provide these. The account of legitimacy she gives (which she calls the 'sovereignty conception') is a form of consent account, which holds that a regime is legitimate insofar as it achieves actual quality consent to rule based on positive governance assessments from a sufficient proportion of its population, where 'quality consent' is consent that meets certain conditions (the details of which I will not go into). On her account, then, for a state to be legitimate with respect to *all* its subjects, it is sufficient for it to have the consent of a certain proportion of its subjects. Her account 'construes legitimacy as a feature of a regime with respect to its subjects *collectively*. To whatever degree a regime is legitimate, its legitimacy applies equally to all the subjects, whether or not each has consented in her individual case'.[112]

Horton does not in fact seem to think that his account has the consequence that existing states are not legitimate (or are only legitimate with respect to some, but not all, of their subjects), suggesting that a different reading is appropriate. I suspect that he in fact holds something similar to Greene, that in order to be legitimate with respect to *all* of its subjects, the state just needs to meet the criteria held to be relevant by *most* of its subjects, or something like that. Greene's motivation for making legitimacy depend on proportional consent is the idea that an account of legitimacy should not lead us to the conclusion that legitimacy is 'virtually unattainable'.[113] Horton similarly contends that we should take the existence of modern states for given and 'seek to explain how political legitimacy is possible'.[114] This may be an appropriate constraint on an account of legitimacy, but it is obviously not an appropriate constraint on an argument against the legitimacy anarchist. In this latter context, it would simply be to beg the question.

[112] Greene (2016) 91
[113] Ibid. 80
[114] Ibid. 139

The claim of my thesis is that the legitimacy anarchist ought to be taken seriously, and we cannot deny this simply by saying that an account of what is necessary to establish a general and exclusive permission to enforce should not have the consequence that no existing states have such a permission (since that is just to assume that the legitimacy anarchist is wrong).

Horton asserts, though, that

> There is … a perfectly straightforward and unobjectionable sense in which we can describe living in a political society as "natural"; and by this I mean no more than that we know of no viable long-term way in which human beings in general can live valuable and worthwhile lives, whatever exactly a worthwhile life is taken to be, outside of a political society.[115]

It is, of course, true in one sense that we 'know' of no such thing: we have not experienced (nor do we have any record of) a 'viable long-term way in which human beings in general can live valuable and worthwhile lives … outside of a political society'. But in the same sense we do not know of a viable way to live in a complex modern society without patriarchy, or in a globalised world without famine, and so on. Just the fact that all societies of the relevant type that we have experience of have been characterised by (for example) patriarchy doesn't mean that this is something we must take as given and that is not apt for our moral criticism. If instead we should understand 'we know of no viable …' as meaning 'we cannot imagine human beings living long-term valuable and worthwhile lives outside of political society' the claim must surely be false. It is possible to imagine a world in which, for example, people coordinate their action by a set of rules that they observe spontaneously without the coercion of any state (and in the strict, metaphysical, sense at least, such a world is not impossible). We may not know how to bring about such a world, but that is a different question. We similarly do not know how to bring about a world without patriarchy, or famine, but we are in no doubt that these are apt targets of our moral criticism. We cannot simply rule out the possibility of a certain type of practice or institution failing to meet our moral standards on the basis of its ubiquity. What Horton might claim instead is that a functioning social community of the sort that could make human life valuable and worthwhile cannot *feasibly* be achieved without political relations (or without exclusive enforcement). Then, again, we get a *feasibility-based* argument against legitimacy anarchism. This is the sort of argument about which my account of feasibility will raise some questions.

---

[115] Ibid. 140

C. Conclusion

In this chapter, I have argued that the two most prominent families of arguments for state legitimacy can only successfully establish the legitimacy of any *existing* states with the aid of premises about feasibility. One family of arguments (some form of which I think underlies the common implicit assumption that state legitimacy is on safe ground, and other forms of which are found in a number of explicit defences of legitimacy) argues from the idea that the state's exclusive enforcement is necessary for some good. It cannot plausibly be claimed, though, that the goods (or good states of affairs) in question make state enforcement strictly *necessary*, but the arguments can be made by arguing that it is the only *feasible* way of achieving something good, or the best. Another family of arguments is based on the consent of the state's subjects, but these cannot plausibly ground the legitimacy of *actual* states (except perhaps with the aid of another feasibility-based argument).

Given the dependence of these arguments on feasibility, it will be worth having some understanding of that concept, and the account I will give in chapter 5 will cast some doubt on whether these arguments can be taken for granted. However, there is another family of views that I have not yet discussed, which, unlike those above, hold that the legitimacy of a certain kind of state can be established *a priori*, that is, independently of any empirical considerations (and so, a fortiori, independently of any feasibility considerations). These arguments are, in a way, similar to the primary kind discussed in this chapter: they too argue that the state's exclusive enforcement is necessary for something of moral importance. However, the arguments above cannot plausibly be understood as claiming that it is strictly *impossible* for the relevant good to be achieved, or for things to be good in the relevant way, without state enforcement: the claim is better understood as a feasibility claim. These a priori arguments, though, claim that there *is* something of great moral importance that could not *possibly* be achieved without the state's exclusive enforcement, and this can be known a priori (it is not a result of empirical facts about the world we live in). These sorts of argument will be the subject of the next two chapters.

Chapter 3:
The a priori Kantian argument

It is quite natural and common to think that state enforcement is justified and permitted in response to the unfortunate circumstances in which we find ourselves, and the arguments discussed in the previous chapter can be understood along these lines. Human frailty and material scarcity combine to make state enforcement necessary for worthwhile human life in society. James Madison put this thought succinctly: 'if men were angels, no government would be necessary'.[116] However, as already mentioned, there is an opposing tradition that holds that state enforcement is justified and permitted *a priori*. However morally good human beings are, and whatever empirical conditions they find themselves in, the existence of an exclusive enforcer is necessary for something of great moral importance (and thus an exclusive enforcer that meets certain conditions is justified and permitted to exist). The first part of this thesis argues that feasibility considerations are essential to establish state legitimacy. It is thus particularly important to look at these a priori arguments for state legitimacy, since they deny precisely this. In this chapter and the next, I will argue that two different kinds of a priori argument both fail in their aim to establish the legitimacy of a certain kind of state without any feasibility premises.[117] In this chapter, I will look at an argument that Arthur Ripstein offers as an interpretation of Kant's *Doctrine of Right*, and in the next, I will look at two similar arguments, one of which is also offered as an interpretation of the *Doctrine of Right* and the other of which is given in a recent book by Philip Pettit.[118]

---

[116] Madison (1990) 267

[117] There is a kind of argument that comes close to claiming that state enforcement is necessary a priori for something of moral importance (in this case, the resolution of conflicts and competing judgments), but does not quite. This is the argument that even morally perfect people may still need state enforcement, made by Gregory Kavka (1996) and (in passing) David Estlund (2016). They argue that there may still be sources of disagreement and conflict even among a society of morally perfect people: in Estlund's example, it may be morally permissible to go as far as violence to procure the one remaining dose of medicine to save one's parent's life (2016, 308-9). Thus, the thought goes, an exclusive enforcer may be necessary to prevent violent conflict even for morally perfect people. However, as Estlund points out, this is not to claim that the need for state enforcement is a priori. There is a *possible* 'arrangement of moral motives' that resolves these conflicts without state enforcement, and indeed, there is a *possible* world where no such conflicts arise. Moral perfection, though, is not sufficient to guarantee either of these things. This does not show that state enforcement is necessary independent of feasibility considerations, since if it *were* feasible to eliminate all of the sources of conflict, for instance, then state enforcement would clearly not be necessary.

[118] Kant (1999); Ripstein (2009); Pallikkathayil (2017); Pettit (2012).

1. The Kantian argument

On a Kantian view, justification, obligating power and legitimacy all come together from one comprehensive argument. This need not be to say that the Kantian view denies any distinction between the three concepts (although it may be that it does, and it is certainly true that Kant does not attempt clearly to separate them out). It is rather to say just that there is one big argument at the centre of Kantian political philosophy that, if successful, establishes all three. My aim will be to argue that even if we accept the core of the Kantian argument, and that justification and obligation follow from it, legitimacy could only do so with the addition of certain feasibility premises: the Kantian hope for a wholly a priori argument from freedom is not met.

In this chapter, I will address the interpretation of the Kantian argument given by Arthur Ripstein in his recent book *Force and Freedom*.[119] In the next chapter, I will address Japa Pallikkathayil's interpretation of Kant, which understands the argument for legitimacy differently.[120] Both present their arguments as exegeses of Kant's *Doctrine of Right*, so if either is right, I will also be addressing Kant's own argument. I will not attempt to assess the exegetical claims of either. For the purposes of this chapter, I will suppose that Ripstein is right in his exegesis, and I will not distinguish between his and Kant's arguments. I will begin by setting out what I take to be the core of the Kantian argument and I will explain how justification and obligation are supposed to follow from it. In the next section I will consider what resources the argument has to establish legitimacy and I will argue that it requires the claim that enforcement is necessary for a rightful condition, which it cannot establish a priori. In the final section, I will argue that not only does the Kantian argument fail to establish legitimacy a priori, but to establish the legitimacy of actual states it requires, like the arguments discussed in the previous chapter, the kind of feasibility premises that cannot be thought to be *obviously* constraints on moral argument.

2. The core

Kant's argument (as presented by Ripstein) begins from the premise of an innate right to freedom as independence. Right, Kant thinks, is about the *external* formal relations between people's choices. It is about the requirements imposed on us by the external relations of our choices with those of others. It is not concerned with the 'matter' or motive of choice (*what* we choose to do, what ends we choose to

---

[119] Ripstein (2009)
[120] Pallikkathayil (2017)

set), but merely its form: how our external actions impact on the *freedom* of choice of others and so, how the external freedom of all can be made consistent. This gives us (analytically) the Universal Principle of Right (UPR): 'Any action is *right* if it can coexist with everyone's freedom in accordance with a universal law, or if on its maxim the freedom of choice of each can coexist with everyone's freedom in accordance with a universal law'.[121] The idea of freedom in question is that of independence. We are free, in this sense, so long as we are not subject, in our purposive activity, to the choices or ends of others: thus, to be free is to have the capacity to set and pursue one's *own* ends.[122] If others set ends for us, then we are not independent of them in the relevant sense. This innate right to freedom, Kant and Ripstein think, extends automatically to our own bodies, since it is by means of these that we act in the world. The innate right to freedom, however, does not immediately give anyone a right to anything beyond their own body. When we are in *physical* possession of an external object (e.g. when we are gripping it), our innate right forbids others from removing it from us. (By removing the apple from your hand, I would interfere with your purposiveness, and use you for my end.) However, this does not give us any right to external objects beyond when we are in physical possession of them; that is, it is not possible on the basis of innate right alone for external objects to remain *ours* even when they are not in physical contact with us. It is an analytic requirement of a property right, though, Kant thinks, that it extend beyond physical possession: 'something external is mine if I would be wronged by being disturbed in my use of it *even though I am not in possession of it*'.[123] Thus, we cannot get property rights on the basis of innate right alone.

Through an obscure argument, Kant argues that UPR, the external freedom of all, requires a system of private property, exclusive rights to external objects beyond physical possession. Thus, a universal system of freedom demands something that respect for innate right on its own is not able to provide. Our individual freedom requires us to be able to have 'intelligible possession' (that is, possession that persists when the object is not in physical possession) of external objects. However, we do not have any rights to external objects that extend beyond our physical possession simply in virtue of our innate right to freedom consistent with a universal system of freedom. Thus, a full system of freedom consistent with a universal law requires 'acquired rights'.

---

[121] Kant (1999) 6:230 (In referencing Kant, I give page numbers of the standard German edition.)
[122] Ripstein (2009) 15
[123] Kant (1999) 6:249

But how are acquired rights possible? Kant's answer is that acquired right is only possible in a 'rightful condition'. Rights to external objects can be acquired simply by taking control of the objects, but these rights can only be provisional (meaning they cannot be enforceable or obligating) outside of a rightful condition. A rightful condition is one in which there is an *omnilateral* legislative will that makes determinate a system of private acquired right, and so makes possible a situation in which an individual's appropriation of external objects does not subject others to the individual's *unilateral* will. It is not just a contingent empirical fact that a rightful condition is needed in order for there to be a system of acquired property rights; it is a conceptual entailment, the idea of a system of private right without a rightful condition is 'morally incoherent from the standpoint of rights'.[124] Without *omnilateral* law, acquisition of external objects is merely a unilateral imposition and so creates no claim of right, but private property rights are necessary for universal freedom, i.e., for right.

Thus, since UPR is the basis of right and obligation, and since UPR requires there to be a system of acquired property rights, we are under an obligation to create and enter a condition of right. Since a system of equal freedom allows you to coerce others to prevent them from coercing you (restricting your freedom), and since a condition of right is necessary for the achievement of a system of equal freedom, you have the right to coerce others to enter into a rightful condition with you. That is, it is permissible to use coercion to establish a legislative (and also, Kant thinks, executive and judicial) authority. If you made and enforced laws on me unilaterally you would be interfering with my freedom as independence, you would be making yourself my master, by subjecting my capacity to set and pursue my own ends to your unilateral will. However, if laws are given *omnilaterally*, that is, in the name of all, they do not make me subject to the will of another. Ripstein argues that this should be understood in terms of the state *making arrangements* for its citizens. In a rightful condition, public officials can be understood not as acting for their own private purposes, but as being empowered to act for others (for the people as a whole). A unilateral will always has a particular end, while the omnilateral will 'only acts to preserve the formal conditions through which people can rule themselves'.[125] This gives rise to the 'principle of public right', which is the principle that 'the sovereign may not give a

---

[124] Ripstein (2009) 146
[125] Ibid. 196

people laws that it could not give to itself, and its corollary that the people must give to itself laws that are the preconditions of its own continued lawgiving'.[126]

The above is a very brief summary of Ripstein's Kantian argument. We can see straightforwardly how this can be considered a *justification* of the state (or of an omnilateral lawgiver). If a rightful condition must be a condition of living in a state, it is a good thing that there is a state because the existence of a rightful condition is a necessary precondition of a system of universal freedom, which is the basic requirement of right. Furthermore, not only is it morally permissible to create a state, it is morally *obligatory* to do so. It also follows from this argument that people have certain obligations towards rightful states (or lawgivers). As we saw, it is morally obligatory to create and enter a condition of right. Further, the rightful laws made in a condition of right *make determinate* the obligation to respect the rights of others (UPR), and they do this by establishing an omnilateral system of property rights. When there is a condition of right, what is required by UPR is that we obey its laws. (The property laws made by the omnilateral lawgiver do not generate new obligations, but give content to an already existing one.) I will turn below to the question why we might think legitimacy follows from the core of the Kantian argument expounded above.

Kant's argument is one that, if successful, would not make legitimacy (or obligation or justification) depend on any feasibility claims. Kant thinks the whole argument can and must be made a priori, without appeal to empirical or 'material' premises. The whole argument is supposed to follow a priori from the idea that everyone has a right to be her own master, i.e., to freedom.[127] Law and coercion are morally required 'however well disposed and law-abiding human beings might be'.[128] That is, they are not required merely as the best feasible option given the violent and conflictual tendencies of human beings. Property rights in external objects cannot be made binding unilaterally, and so for them to exist there must be an omnilateral lawgiver. This follows conceptually from the idea of right; it is not just empirical facts that make this true. Thus, on the Kantian view, obligation and legitimacy do not depend on the state's being the best feasible alternative, but simply on its securing a rightful condition.

---

[126] Ibid. 25
[127] Ripstein (2009) 5
[128] Kant (1999) 6:312

3. The assurance argument

I think there are good reasons to doubt the success of the core argument described above that I do not have space to go into here,[129] but I want to set aside these doubts and suggest that even if we accept the core of the Kantian argument and we accept that states' obligating power follows from it, the exclusive right to enforce does not. What we have seen so far is an argument for the claim that the imposition of a single set of binding laws is necessary for the existence of a rightful condition, and so it is obligatory to create one. All this gives us is the claim that it is permissible for the state (or some similar agency) to create binding laws, i.e., to create a rightful condition, since if it is obligatory to create one it must be permissible to do so. (This is really a combination of the permission to create laws and the power to make laws that are binding.) However, it does not follow from this that the state is permitted to *enforce* those laws. Yet it seems that Ripstein and Kant think not only that a system of omnilateral law is necessary for right, but also that an exclusive enforcer (a state) is too. Ripstein talks of the necessity of a system of 'public enforcement' and Kant of the necessity of a 'common and powerful will'.[130] It sounds as though what they are talking about is a single, exclusive enforcer (perhaps identical to the omnilateral lawgiver). How, then, does the state get this permission to enforce, i.e., legitimacy?

I think that the beginnings of a Kantian answer come simply from the claim that the authorisation to coerce is just the flipside of a right. Kant says that coercion that is 'a *hindrance of a hindrance to freedom*' 'is consistent with freedom in accordance with universal laws, that is, it is right. Hence there is connected with right by the principle of contradiction an authorisation to coerce someone who infringes upon it'.[131] A right is what is allowed by freedom in accordance with universal laws. Anything that is 'a hindrance of a hindrance to freedom', Kant thinks, is therefore allowed by right. Any interference with right, then, can be coercively prevented by right.

So, if the state establishes a system of law that is a system of public right (that is, it establishes a system of property rights that creates a rightful condition), then it is automatically permissible for the state to enforce these laws, since by doing so

---

[129] The core argument is only as strong as the argument that a system of property rights is necessary for universal freedom. A.J. Julius (2017) powerfully argues that this latter argument is unsuccessful. If he is right, no more would be necessary to show that the Kantian argument cannot establish legitimacy, but I do not have space to go into Julius's reasoning here, and so I will set aside such doubts. There are other powerful challenges to the core in Sangiovanni (2012) and Valentini (2012c).
[130] Ripstein (2009) 159 and Kant (1999) 6:256
[131] Kant (1999) 6:231

they would merely be enforcing right and right is necessarily enforceable. However, it is permissible for *anyone* to hinder hindrances to freedom, to coercively enforce right. The question arises, then, how the state gets the *exclusive* right to enforce its laws. That is, how does the state get the right to be the *only* enforcer of public right and to coercively prevent others from using coercion to do the same? Whatever gives the state the exclusive right to create authoritative laws does not automatically also give it the *exclusive* right to coerce.

One answer we might be tempted to give is that an exclusive enforcer is necessary for the system of laws to be effective. The state is permitted to set itself up as an exclusive enforcer because it needs to do so in order for its system of laws to be adhered to and so for rights to be respected. However, this cannot be the Kantian answer because it depends on a contingent empirical feasibility fact. It is not true *a priori* that a system of laws can only be effective when enforced by a single enforcer. There is no (non-empirical) reason that there could not be a society of angels who obey a single dominant system of rules without their doing so being enforced, or a society where multiple enforcers all agree to (and do) enforce a single system of rules without any higher enforcer ensuring that they do. If you think law is necessarily coercive, then you will not want to call such a system of rules a system of law, but it could be just as effective at creating a condition of right if universally obeyed.

I think the answer to this problem that we find in Ripstein (and, by extension, if he is right, Kant) is that rights are not in fact enforceable (and therefore there are no rights) unless we have public assurance of the compliance of others.[132] If it is true that assurance is necessary for there to be enforceable rights, and enforcement is necessary for assurance, then we have an argument for the necessity of the enforcement of rights for a rightful condition. This is not quite yet what we need. It might be thought, in addition, though, that a *single exclusive* enforcer of rights is necessary for there to be public assurance of compliance. If enforcement is necessary for assurance at all then this is presumably because, when a system of rules is set up, the only way that we can be sure that it will be universally complied with is if compliance is enforced. If there are multiple enforcers, the situation is no different to how it would be if there were no enforcers. If we cannot have assurance that people will comply with a single system of rules with no enforcement, then we similarly cannot have assurance that multiple enforcers will all comply with (enforce) a single set of rules. Thus, if

---

[132] Ripstein (2009) 159ff.

enforcement is necessary at all, a single supreme enforcer who can enforce the compliance of any other secondary enforcers will be necessary.

Why, though, think that assurance is necessary for a condition of right in the first place? Ripstein's answer is that there in fact is no obligation, and so no right, if there is no assurance of the compliance of others. The duties imposed by right are conditional on the conduct of others because 'all obligations of right must be within a system of right "in accordance with universal law"'.[133] If you have no assurance that I will comply with the system of right and so respect your property, you have no obligation to respect my property and so I cannot coerce you to respect my property: 'if either of us refrains from taking what belongs to the other without assurance, we restrict our choice on the basis of the other's particular choice, rather than in accordance with a universal law'.[134] In general if people have no assurance, they have no obligation and so cannot permissibly be coerced; we do not have a condition of right. Thus, if the existence of a single supreme enforcer (a state) is a necessary part of there being assurance, and there being assurance is a necessary part of a condition of right, then the state must have an exclusive right to enforce its laws, since there being an exclusive enforcer is an essential part of a system of universal freedom.

Let us grant (for the sake of argument) that the above arguments are sound. Suppose that assurance is a necessary part of a condition of right and suppose that *if* enforcement is necessary for assurance, then a single exclusive enforcer is necessary for assurance. However, again Kant needs the premise that enforcement is necessary for a condition of right or, more specifically, for assurance. What I cannot see is that Kant is entitled to this claim at all, without bringing in contingent empirical feasibility claims. How is it that coercive enforcement provides assurance in the first place? The answer must be that it creates certain *expectations* about how people will behave. Given what we know about human preferences, when the state coercively enforces certain property rights, we come to *expect* most people not to interfere with those rights because we expect them to prefer avoidance of the punishment to whatever benefit could be gained by infringing the right.[135] With the exception of certain special cases, coercive

---

[133] Ibid. 160
[134] Ibid. 162
[135] In some cases, coercive enforcement might consist not in threatening punishment, but in making infringement physically impossible (or hard). This is not the usual case, though, and there is only so much that the state can make impossible. It is hard to imagine a state whose system of laws is enforced entirely by making infringement impossible. If the Kantian argument could only ground this sort of enforcement, it would fail to ground most of what states actually do.

enforcement does not *guarantee* compliance; there is no entailment from one to the other.

With this in mind, it seems that there is no a priori reason that we cannot have just the same sort of assurance of respect for a system of rights without enforcement; in other words, enforcement is not *a priori* necessary for assurance. There is no a priori reason why it would not be possible for there to be a society of angels in which everyone accepts and observes a single set of laws without any enforcement. If this was a general practice, we would come to *expect* people to comply with the laws in just the same way we do as a result of enforcement. If enforcement is not a priori necessary for assurance at all, then a single exclusive enforcer cannot be necessary for assurance. (The argument above that a single exclusive enforcer is necessary was premised on enforcement being necessary in the first place for us to be sure of general compliance, but we do not have any a priori reason to rule out the possibility of this sort of assurance without *any* enforcement whatsoever, let alone without a single exclusive enforcer.) The only argument that there seems to be in Ripstein, however, that might establish an *exclusive* permission for a state to enforce its laws and to determine who may wield enforcement power in its territory, is this argument under discussion from the necessity of assurance for a condition of right. If a single exclusive enforcer (or enforcement at all) is not necessary for assurance, then we have no reason to think that a single exclusive enforcer is necessary for a condition of right, and no reason to think that right requires rightful states (omnilateral lawgivers) to have an exclusive permission to enforce.

Above I accepted without question Ripstein's argument that assurance is necessary for a condition of right because we can have no obligation to respect the property of others without assurance that they will also respect ours. In fact, however, I have doubts about this argument. If we imagine a situation with universal laws legislated by some public omnilateral lawgiving authority but no form of assurance of the compliance of others, then it seems to me that if anybody in such a situation restricts our choice on the basis of these laws, it is just as much in accordance with a universal law as it would be if we had assurance. If you do not comply with the law, but I do, in what sense am I restricting my choice on the basis of your particular choice? Japa Pallikkathayil expresses similar doubts, asking how respecting your rights in the absence of assurance can permit you to violate mine.[136] She, however, does not deny that there is an assurance problem in

---

[136] Pallikkathayil (2017) 41

the state of nature, making a state with enforcement power necessary for a condition of right. Instead, she proposes a slightly different understanding of the assurance problem. The problem is not that one cannot have enforceable rights without assurance. (Of course, one cannot have enforceable determinate property rights, on the Kantian view, without an omnilateral lawgiver, but this problem only establishes the need for the lawgiver, not for an enforcer, so we need to imagine a scenario in which you have determinate property rights but lack the assurance that they will be respected.) Rather, she says, the problem is that whatever assurance you can have in the state of nature that one's property rights will be respected must rely on the unilateral will of others.[137] So long as my assurance that others will respect my property rights relies on the unilateral decisions of others to respect them, I am still dependent on their choices, and so not free. We do not have a condition of right.

If the assurance problem is to be understood as Pallikkathayil suggests, then the argument I gave above against Ripstein loses traction. Even in a society of angels, where we can have extremely good predictive assurance that others will comply with a system of rules, we are still dependent on the unilateral choices of these others, since they *could*, counterfactually, choose *not* to comply. There is nothing to stop them. The same is true in a society of multiple enforcers. If we understand the assurance problem in this way, we get a different argument for the necessity of an exclusive enforcer for universal freedom. Without one, the argument goes, however angelic our fellows *in fact* are, we will always be dependent on their choices, since they have the *capacity* to subject us to their ends. I will address this argument and one by Philip Pettit which it very closely resembles in the next chapter. The lesson for now will be that Ripstein has no apparent *a priori* argument for a state's *exclusive* permission to enforce, *unless* he turns to this counterfactual sort of argument.

4. What feasibility premise?

I have shown that Ripstein does not succeed in his aim of providing an a priori argument for state legitimacy. The argument he gives does not plausibly succeed without the empirical claim that it is not feasible to achieve assurance of the compliance of others without an exclusive enforcer. There is no *a priori* reason why there could not be just the same sort of assurance without an exclusive enforcer, but it may well be *infeasible* (empirically) to achieve this. Since it is such a *feasibility* claim that most plausibly would allow this argument to go through, I

---

[137] Ibid. 38, 42

think this also supports my claim that plausible arguments for state legitimacy depend on feasibility premises. However, it could be thought that it is so robustly infeasible to achieve assurance without an exclusive enforcer that the above should be treated as, for all intents and purposes, an *impossibility* claim. In chapter 2 we looked at an argument that depended on the claim that it is not feasible to establish a dominant set of rules in a society that are universally adhered to except by exclusive enforcement. There, similarly, we considered the objection that worlds in which a dominant set of rules are established in a society without exclusive enforcement are so distant that this claim should be treated as an impossibility claim, as a straightforward constraint on moral requirement: the worlds in which it is not true are just too distant to be relevant to moral permissibility. There, I accepted this for the sake of argument and argued that the argument's success relies also on feasibility premises that cannot plausibly be thought to be of this sort.

The same is true here. The success of Ripstein's argument in establishing legitimacy for at least some actual states depends not only on the premise that it is not feasible to achieve assurance of the compliance of others without an exclusive enforcer, but also on the premise that it is not feasible to achieve assurance without the exclusive enforcement of *the particular state in question*. We noted in the introduction the importance of particularising the permission to enforce. It will not do to show that *some* state is permitted to be an exclusive enforcer, since this will only show that it is *possible* for states to be legitimate, not that any actual states are. Rather, we need to be able to show that *some particular* state is permitted to be an exclusive enforcer. In order to do this, we would need to show that this particular state's exclusive enforcement is necessary for assurance (or it is not feasible without it).

In general, that fulfilling some task is necessary to solve some weighty moral problem is not, intuitively at least, enough to show that it is permissible for any particular agent to fulfil that task. It may be that there are many agents capable of fulfilling a task, and it could be that some of them are not permitted to fulfil the task precisely because their doing so would prevent others from fulfilling the task. Some of the capable agents could have a claim not to be prevented from fulfilling the task because of some personal connection of theirs to it, or because they would fulfil the task *better*. For example, suppose that I am in a bad mood and there are a number of people in a position to lift my spirits by taking me out to dinner, including Alex and Brian. Despite this, it could be that Alex would in fact be doing something *wrong* by taking me out to dinner, because he would thereby be

preventing Brian from doing so. Perhaps Brian has a special claim to fulfil this task because he is my partner, or perhaps Brian would simply be *better* at lifting my spirits. Analogously, I think, the fact that a state is fulfilling some necessary task for the society over which it governs and that its acting as supreme enforcer is necessary to the fulfilment of this task is not yet enough to show that it has an exclusive right to enforce. By doing so it is preventing other agents or institutions from fulfilling the task in question. It could be that some *other* agent or institution has a special claim on this task or that some other agent or institution would perform this task much better.

The argument for the permissibility of enforcing rightful laws does not simply appeal to the necessity of this task, but rather to the analytic enforceability of right. The fact that a state establishes a *rightful* condition is enough to show that its laws can be enforced. However, the argument for the state's *exclusive* permission to enforce seems to need to appeal to the necessity of exclusive enforcement for a particular task: achieving assurance. But it might be that another agent or institution has a special claim to establishing assurance of a rightful condition or would do so better. Thus, the feasibility premise that Ripstein's argument would need is the strong one that the exclusive enforcement of this particular state (the one for which we are trying to establish legitimacy) is the only feasible way of achieving assurance of a rightful condition in the given circumstances. Again, (as with the arguments discussed in the previous chapter), feasibility premises are needed both to show that *the* state (some state) is necessary for right, and to show that any *particular* state is necessary. Again, the feasibility premises required for the latter are less obviously strict constraints on moral argument in general.

Chapter 4:
The a priori counterfactual argument

In the previous chapter, I argued that one notable attempt to defend the legitimacy of a certain kind of state a priori (Ripstein's Kantian argument) fails in its aim. There is no *a priori* reason to suppose that assurance is only possible with the state's exclusive enforcement. However, we noted that there is a different a priori argument available, which Pallikkathayil offers as an alternative interpretation of Kant. On this view, the problem with the stateless society is not that we lack assurance of the compliance of others, but rather that we are counterfactually dependent on the choices of others, since they *could* choose not to respect our rights if they wished to, and this is supposed to be an a priori necessary feature of stateless society. Since freedom on the Kantian view is *independence* of the choices of others, this means that unfreedom necessarily characterises society without state enforcement.

Philip Pettit, in his book *On the People's Terms*, offers a similar (although not identical) argument for the necessity of state enforcement for freedom a priori. Pettit takes the problem of political legitimacy (which I think he understands to include legitimacy in my sense) to be that of showing that state enforcement (and, presumably, the existence of an exclusive enforcer) can be consistent with individual freedom. In his book, Pettit makes an empirically informed argument that state enforcement *can* be consistent with freedom, but he also gives an argument which purports to show, like the Kantian argument, that state enforcement is a priori *necessary* for freedom. If this can be shown, and if respecting freedom is morally required, then state enforcement will be morally obligatory a priori (for those in a position to bring it about). And if it is morally obligatory, then, a fortiori, it is morally permissible (and if what is necessary is an exclusive enforcer, then it is presumably also morally permissible for the state to be an exclusive enforcer). (We might question the claim that demonstrating legitimacy just amounts to demonstrating the compatibility of state enforcement and *freedom*, or that respecting freedom is morally required, but it seems at the very least that a significant step will have been made towards demonstrating legitimacy if this can be shown.)

Like Pallikkathayil's, Pettit's argument appeals to the idea that some unfreedom is ineliminable from even an ideal stateless society because of the counterfactual

capacity of people to interfere in the affairs of others. As mentioned, most of Pettit's extensive and interesting discussion of state legitimacy does not purport to be a priori. The argument I am interested in occurs quite briefly towards the end of the chapter.[138] It is presented by Pettit as showing that the democratic state is the source of something *good* and so is something to be welcomed, not merely a remedy for unfortunate (if unavoidable) circumstances. The idea is to argue that even in an *ideal* stateless society (a Kantian kingdom of ends), the introduction of a state (an exclusive enforcer) would be an improvement because it provides something good that is unavailable without it. The good in question is the good of 'status freedom'.

One is only truly free, on Pettit's view, when one has this 'status freedom', that is, when one has the status of an independent member of society. To be free for Pettit is not to be *dominated* by any other will, where a person, A, is dominated by an agent or agency, B, 'to the extent that B has a power of interfering in [A's choice] that is not itself controlled by A'.[139] What matters for A's freedom as non-domination is not whether anyone in fact does interfere in her choices, but whether anyone has an uncontrolled *power* to interfere in her choices. Further, says Pettit, 'to enjoy the relevant freedom of non-domination is to be someone who commands a certain standing amongst your fellows'.[140] To have this kind of standing, this equal status, is not to be subject to the uncontrolled power of interference of any other.

Pettit's argument, then, is that full status freedom would not be achieved in an ideal stateless society in a way that it could be achieved with the addition of a state. The ideal stateless society, or kingdom of ends, is one populated by morally motivated people who, Pettit assumes, are committed to showing respect and concern for others.[141] This, he says, commits them to not interfering in the basic liberties of others and to providing resources required by others' basic liberties. However, even despite this, people's status freedom would not be fully assured in this ideal stateless society:

> Were people moral in the degree imagined, then the more wealthy and powerful would be disposed not to interfere with others and not to allow others to go

---

[138] Ibid. 181
[139] Ibid. 50
[140] Ibid. 91
[141] Pettit's argument does not depend on the persistence of conflict among the perfectly morally motivated (see footnote 117), so he could also assume that the ideal stateless society is one where conflict-generating scarcity is absent. (His conclusion, then, is not merely that even the perfectly morally motivated need state enforcement, but the even stronger claim that state enforcement is necessary for freedom *a priori*).

> without needed resources. But their acting on that disposition would depend on their continuing to embrace and abide by the requirements of the assumed morality. It would depend … on their displaying a good will rather than a weak will or a will to evil. In such a world, therefore, some members would have to depend on the goodwill of others for enjoying the basic liberties.[142]

Thus, the addition of state enforcement could provide something that could not be provided in an ideal stateless society: it could make each member of society independent of the will of others. The threat of enforcement, even if not necessary to ensure compliance with just laws, *guarantees* this compliance and removes *choice* about whether to show respect and concern for others. Thus, the addition of a state need not change anything about people's actual behaviour, nor need its enforcement involve any actual use of force, but the counterfactual enforcement that it establishes gives people a kind of status freedom that they could not have before. Thus, if a particular state protects the status freedom, or freedom as non-domination, of its subjects (and perhaps if it is the only one that could do so in the circumstances), then it might be thought that it will be legitimate.

Both Pettit and Pallikkathayil, then, need to argue that there is a problematic sort of unfreedom (dependence on others) that persists even in an ideal stateless society (a society of angels, a kingdom of ends) and that can be alleviated by the introduction of a state. I will first take note of, and endorse for the most part, some arguments recently made by Thomas Simpson. They do not, though, show, as they purport to, that republican freedom is *impossible*, if we understand it as *status* freedom. I then argue that given this understanding of republican freedom, and given an assumption of Hobbesian rough natural equality, Pettit is wrong to claim that there would be unfreedom in an ideal stateless society made up of individuals. However, I then note, if we admit group agents into the ideal stateless society (and there is no reason not to), there may be troubling unfreedom in the ideal stateless society. But, I argue, the sort of unfreedom that would exist in an ideal stateless society with group agents cannot be alleviated by the introduction of a state. I then relax the assumption of Hobbesian natural equality and show that, although this does have the consequence that there is status unfreedom even in the ideal stateless society, again the state can do nothing to alleviate this. Finally, I consider a view that takes the sort of freedom that matters not to be status freedom, but independence of the unilateral will of others (Pallikkathayil's view). I argue that still the state is no improvement on the ideal stateless society in terms of freedom understood thus.

---

[142] Ibid. 182

1. The possibility of republican freedom

Thomas Simpson has recently argued that there will *always* be domination of the sort that Pettit is concerned by, and so that republican freedom is *impossible*.[143] He argues that it is possible to be dominated by an uncoordinated collection of agents who have the uncontrolled *power* to coordinate to interfere (without forming a group agent). However, nobody can protect you from such power except other people, who will also have the uncontrolled capacity to coordinate to interfere, and so dominate you. As Pettit himself thinks, although the state might protect against the sort of domination that there would be in a state of nature, its subjects need to be protected from it (so as not to be dominated by it). The only way they can be (which Pettit points to) is by the capacity of the people to resist the state. But if the people have this capacity, they must also have the uncontrolled power to interfere with any individual subject of the state. Thus, although Pettit is right (on Simpson's view) that there will always be domination in a stateless society, no matter how ideal, what this misses is that there will always be domination in a state society as well. This, if right, provides an easy way to reject Pettit's a priori argument. Though one is inevitably dominated in an ideal stateless society (by other individuals and potential collections of individuals), one is also inevitably dominated in a state society (either by the state itself or by potential collections of other individual subjects of the state) and there is no reason to suppose that the domination in a state society is less significant or somehow preferable from the point of view of freedom to the domination that there would be in a stateless society. I think that Simpson's argument is right almost all of the way, but I do not think it quite establishes its conclusion. I will thus briefly run through what I think Simpson gets right, and then explain why his argument will not do to answer Pettit's a priori argument.

Simpson's argument is centred around two cases:

> *Nearly Coordinated Masters*: There are three Masters and a Slave. No Master is strong enough alone to interfere with the Slave. If two of them coordinate, they can interfere. Master$_1$ is ready to interfere, but Master$_2$ and Master$_3$ are benevolent, rejecting the proposed joint intention and not engaging in team reasoning.

> *Non-Coordinating Masters*: As *Nearly Coordinated*, but all three Masters are benevolent.[144]

---

[143] Simpson (2017). Similar arguments seem like they might well apply to the Kantian sort of freedom as independence that Pallikkathayil is interested in.
[144] Ibid. 32

On Pettit's account, the slave in *Nearly* is dominated by Master$_2$ and Master$_3$, since either has the capacity to bring about interference with the slave at will. Simpson argues that the slave in *Non* is also dominated. In both cases, ingratiating himself with the masters is an appropriate strategy for the slave. In *Non*, two masters would need to turn on the slave, not one, but this does not seem an important difference: the slave remains dependent on the good graces of the masters.

Pettit, however, might just claim that domination must be a dyadic relation, and so there can be no domination in *Non*. For this reason, Simpson argues that *polyadic* domination is possible. First, he argues that it is possible for collections of individuals (groups that do not constitute group agents) to have the *power* to interfere jointly. That such groups can have the power to act jointly is illustrated by the case of 'holidaymakers on a beach, who form a chain to pull a swimmer out of difficulty'.[145] It would be wrong to say that they do not have the power to rescue the swimmer: they have this power *together*, 'even though − because the chain is not yet formed − no individual does'.[146] Further, he argues, when a collection of individuals has the joint power to interfere, it is implausible to say that it could not be dominating: 'Interference by a group is just as worrisome, perhaps more so, as interference by an individual. A group of agents may be less likely to interfere, but it is unclear why their collective capacity to do so at will should not compromise my freedom in the same way that an individual's capacity does'.[147] We have just the same sorts of reasons to ingratiate ourselves with the group that has the collective power to interfere, and so on. I find these arguments persuasive: there seems to be no good reason to rule out the possibility of *polyadic* domination. And thus, since if groups of citizens of a state have the power to interfere with the state, they must also have the power to interfere with individuals, it follows that in a state that is not *itself* dominating, individuals will be polyadically dominated by potential groups of coordinated fellow citizens.

I think that Simpson is right, then, that domination, conceived simply as the capacity to interfere at will (either individually or jointly with others) is unavoidable. However, I think that some of what Pettit says suggests that the kind of instances of domination that are morally problematic, or that we ought to be concerned to eliminate, are only a subset of the instances of domination tout court. Thus, I do not think we yet have a reply to the a priori argument, since it

---

[145] Ibid. 37
[146] Ibid. 38
[147] Ibid. 42

might be that the *problematic* kind of domination is ineliminable from the ideal stateless society, but can be remedied by the introduction of a state. We can see this way out if we take seriously Pettit's claim, mentioned above, that 'to enjoy the relevant freedom of non-domination is to be someone who commands a certain standing amongst your fellows'.[148] This is what is meant by 'status freedom'. The idea, I suppose, is that the status freedom that matters is freedom from subjection to *asymmetrical* domination relations with others. When whatever domination one suffers at the hands of others is perfectly matched by exactly equivalent domination they suffer at your hands, it seems fair to say that one's 'standing' or 'status' is that of a free person.

Free persons, those who have status freedom, says Pettit, 'can walk tall, and look others in the eye. They do not depend on anyone's grace or favour for being able to choose their mode of life'.[149] This leads to Pettit's proposal of the 'eyeball test' for determining how far the basic liberties should be safeguarded: they should be protected up to the level where all 'can look others in the eye without reason for the fear or deference that a power of interference might inspire; they can walk tall and assume the public status, objective and subjective, of being equal in this regard with the best'.[150] It seems clear that this eyeball test will be met, and all members of society will have status freedom in as full a sense as possible when all domination or dependence on the good-will of others is fully symmetrical. If, on the other hand, you are dominated by another in a way that you cannot match, your status is in some sense *inferior* to theirs and the eyeball test may not be met. This, it might be thought, is the true or troubling kind of unfreedom. It is in this sense that status freedom is bound up with status equality. For all to have status freedom is for there to be status equality. It is only possible for nobody to have inferior status if all have equal status.

If what matters is *status* freedom, and thus status *equality* (the absence of *asymmetrical* domination), then it may still be that republican freedom (of the sort that matters) *is* possible. If the only sort of domination present is domination by potential coordinated groups among, say, the people who keep checks on the power of the state, there will only be troubling (asymmetrical) domination if there are differences in people's capacity to interfere with others. If all individuals had equivalent capacities (which is not impossible) domination would be fully symmetrical. Though of course A would be dominated polyadically by the

---

[148] Pettit (2012) 91
[149] Ibid. 82
[150] Ibid. 84

potential group of B, C and D in a way that the potential group of B, C and D would not be dominated by A, A's dependence on each individual as part of a potential group would be exactly matched by that individual's dependence on A as part of an equivalent potential group involving A. It might be thought, then, that the trouble with the ideal stateless society is that it involves *asymmetrical* domination (and Pettit points to the dependence of those with less wealth and power on those with more), and that this *asymmetrical* domination could be removed by the introduction of a state. If the state's capacity to interfere were fully controlled by its citizen body, then it could succeed in removing asymmetrical domination, since the polyadic domination *among* members of the citizen body need not be asymmetrical. Let us see, then, whether the case can be made that there would be asymmetrical domination in the ideal stateless society of the sort that a state could eliminate.

2. Mutual domination

There are two ways of interpreting the ideal stateless society as conceived by Pettit. We have said that the inhabitants of the ideal stateless society are morally motivated and committed to showing concern and respect for their fellow inhabitants. Either these inhabitants of the ideal stateless society agree on and abide by a common set of rules or they do not. Let us start with the first possibility.

Part of the plausibility of Pettit's claim that some would be dominated by others or depend on the good-will of others rests on there being differences of wealth and power. Those with less wealth and power will depend on the good-will of those with more, even if the wealthy and powerful are perfectly disposed never to interfere in the choices of the less wealthy and powerful. If people in the ideal stateless society agree on and abide by a common set of rules, then we can imagine these rules including a set of property rules, and so we can easily make sense of such differences in wealth and power. However, if avoiding domination is a requirement of justice or morality (and if it is not, there is no complaint against the ideal stateless society), then ideal morally motivated agents ought to agree on and abide by a set of rules that, so far as it is possible for a set of rules to do so, rules out relations of domination. Thus, if A's having more wealth or power than B allows A to dominate B, by making B dependent on A's good-will, the inhabitants of the ideal stateless society ought to agree on rules that rule out inequalities of wealth and power sufficient to make domination possible. There presumably are possible sets of rules that a society of perfectly morally motivated

individuals could agree on and keep to that would restrict inequalities of power and wealth. Thus, in an ideal stateless society with a common set of rules, there would be no domination arising from differences of wealth or power. Further, if we make a Hobbesian assumption that people's *physical* strength or power is roughly equal (an assumption that I will relax later), there would be no domination of the sort Pettit envisages at all.[151]

However, it might be objected that there still is domination, since even if all members of this ideal stateless society abide by rules preventing inequalities of wealth and power, they still could counterfactually interfere in the choices of others if they wanted to. On the one hand, people *could* accumulate wealth or power if they broke the collectively agreed-on rules (so long as the rules that make the social wealth and power possible are still followed). On the other hand, even without wealth or power, people still *could* interfere in the choices of others. In both of these respects, though, whatever domination there is (or counterfactual capacity for interference) is symmetrical for all (assuming still that we have Hobbesian rough natural equality). All are in exactly the same position with regard to the possibility of accumulating wealth or power, or the possibility of interfering in the affairs of others. If you and I are both members of the ideal stateless society, you *could* counterfactually accumulate wealth or power and so dominate me (and so, it might be thought, because domination is a counterfactual capacity, you already do dominate me now), but equally I could counterfactually accumulate wealth or power and come to dominate you. Thus, we both dominate each other equally (and, indeed, in exactly the same way). In fact, everybody in the ideal stateless society dominates every other equally in these respects. Domination is fully mutual. There might, further, be *polyadic* domination, of each member of society by potential coordinating groups. But, as argued above, where the power of individuals is equal, this polyadic domination will be symmetrical for all. There thus *is* domination in the ideal stateless society (with Hobbesian rough equality), but not asymmetrical domination, and so seemingly not status unfreedom. In other words, the domination there would be does not seem to be of the kind that Pettit considers troubling (or indeed, given Simpson's arguments, of a kind that is not inevitable).

It might be thought that even mutual domination is morally troubling (as, I think, Pallikkathayil's Kant will say). If we take this route, as we have seen, troubling unfreedom will be inevitable. Nevertheless, I will return to this view below to see

---

[151] Hobbes (1996) 82 (13:1-2:60-1)

whether there is any reason to think that the state could improve on the domination or dependence inherent in the ideal stateless society.

I have argued that if the inhabitants of the ideal stateless society agree on and abide by a common set of rules there will be no status unfreedom. The other possibility is that the inhabitants of the ideal stateless society do not agree on a common set of rules. However, property accumulation and powers and privileges associated with offices or positions can only make sense against the background of a set of rules defining property rights and conferring powers on individuals. Without common agreed rules, then, there can be no inequalities in wealth or power beyond what results from natural inequalities in strength or ingenuity. However, if we assume Hobbesian rough natural equality, then there can be no significant inequalities in these. Thus, as above, if there are no agreed-on rules in the ideal stateless society, the capacity of all members of society to interfere with the choices of others will be roughly equivalent. Thus, here too, whatever domination there is will be symmetrical.[152]

If we assume rough natural equality, then, there is no status unfreedom in the ideal stateless society. Thus, it is possible for there to be status freedom without state enforcement: it is not an a priori requirement of freedom as such.

3. Collective domination

It might be said that my argument thus far has been too quick. The argument above needed two assumptions to show that all domination in the ideal stateless society is symmetrical: that the only agents capable of domination in the ideal society are individuals, and that there is rough natural equality. If either of these assumptions does *not* hold, it might be thought, there *would* be asymmetrical domination (and so status unfreedom) in the ideal stateless society. Since an ideal stateless society without group agents and natural inequalities is *possible*, it is still the case that state enforcement is not an a priori requirement of freedom *as such*. But it might be argued that state enforcement, though not required a priori by freedom *as such*, is required a priori for freedom in societies with group agents or natural inequality (which may be ordinary features of human society).[153] I will

---

[152] My argument in this section, then, shows that Simpson's claim that republican freedom is impossible is not quite correct (at least if we understand republican freedom as status freedom, as Pettit does). At least in a society characterised by equality (natural equality plus property rules that rule out inequalities in wealth or socially constituted power of the sort that give rise to status unfreedom), there will be no asymmetrical domination, no status unfreedom, and so no troubling republican unfreedom.

[153] It might then be that an argument to the permissibility of state enforcement (for a certain kind of state) is available that does not depend on feasibility premises: natural

take each in turn and argue that, although relaxing these assumptions has the consequence that there *is* troubling unfreedom in the ideal stateless society, the introduction of a state would be no improvement.

First, then, it might be thought that individuals are not the only sorts of agent dependence on whose will could amount to domination. (Potential coordinated groups are not agents, even when coordinating.) As well as individuals, there may also be *group* agents, corporate groups with a structure giving the group the sort of unity that allows it to have and express attitudes, to set and pursue purposes and so on.[154] (Some may want to deny the possibility of such group agency, but there is no need for me to do so.)

There is not obviously anything inherently morally troubling about the existence of corporate agents (and it might even be thought that the opportunity to become a member of such a corporate agent is valuable for individuals). Thus, there is every reason to think that the existence of group agents is consistent with the ideal stateless society. And because group agents are capable of setting and pursuing their own purposes, because, in short, they have a *will*, it is possible for them to dominate other agents (whether groups or individuals).

In an ideal stateless society in which group agents are present, some agents will be dependent on the good-will of others in a way that is *not* entirely mutual. Even if all *individuals* are roughly equal in natural power, *groups* of individuals will obviously be more powerful collectively than any individual (and may be more or less powerful than each other). As we saw before, agents in the ideal stateless society are dependent on the good-will of others, since others *could* interfere with their choices (even though we know they in fact will not). I argued that this is not problematic because fully mutual. However, the dependence of an individual on a corporate agent is not reciprocated. The corporate agent, by virtue of being made up of multiple individuals, will be more powerful than the individual, so the individual will be dependent on the collective in a way the collective is *not* dependent on the individual. The mere existence of corporate agents, then, creates domination. Even if it will not be exercised, a group agent automatically has a capacity to interfere in the choices of individuals which those individuals have no recourse against; the individuals have no equivalent capacity.

---

inequality might be strictly *impossible* to eliminate, and group agents, though presumably possible to eliminate, might be thought *desirable*.

[154] I will use 'group agent', 'collective agent' and 'corporate agent' interchangeably.

Let us grant that this sort of non-mutual domination is troubling and can count as a lack of status freedom; the individuals who are dependent on the good-will of corporate agents fail the eyeball test: they could not look the corporate agents in the eye (metaphorically speaking) without reason for fear or deference. Someone's not being a member of a group that dominates them might constitute a troubling lack of status. However, I do not think that the introduction of the state can eliminate non-mutual dependence on the will of corporate agents, and so the presence of this sort of dependence in even the ideal stateless society is no argument for the necessity of state enforcement for freedom.

The state can eliminate *some* dependence on the will of corporate agents. It may be that the state can use its enforcement power to remove the uncontrolled capacity to interfere of all group agents existing within its jurisdiction. It cannot make group agents' interference any more *predictively* unlikely, since in the ideal stateless society it is already the case that all agents are perfectly morally motivated, and so they will not in fact interfere against another agent's will. However, the state can either *control* their capacity to interfere, making it non-dominating, or remove the capacity altogether. It can make it the case that *if*, counterfactually, a group agent decided to interfere with another agent, it *could* not (or would find it costly to do so).

However, there remains one group agent whose capacity to interfere the state cannot control. That is the state itself. Pettit is committed to the state being an agent.[155] Indeed, for anyone who holds that the sort of freedom that matters is independence of the *will* of others, it will not be a plausible line of defence to maintain that the state is not an agent and so not capable of dominating because it has no will. This would have the consequence that *no* properly constituted state could possibly dominate or limit the freedom of its subjects (or others), which surely is unacceptable. Thus, let us assume that the state's superior capacity to interfere in the choices of its citizens, which it must have if it possesses exclusive enforcement power, is potentially dominating. Even if the state is perfectly well-intentioned, it *could* interfere at will in the choices of its subjects. Unless, that is, this capacity is suitably *controlled* by its citizens. If it *is* possible for the citizens to achieve suitable control of the state (and so, for the state not to be able to interfere at will), then asymmetrical domination could be eliminated, since the kind of domination that would remain among the citizen body (highlighted by Simpson) could be *symmetrical*. The problem of achieving suitable control over the state,

---

[155] Pettit (2012) 133

Pettit says, is the problem of legitimacy. If citizens have *individualised, unconditioned* and *efficacious* control of the interference power of the state, he argues, then it does not dominate them. It is, he maintains, *possible* for a democratic state to be suitably controlled in this way.

The requirement of *unconditioned* control is particularly important for the purpose of marking the difference between the ideal stateless society and the democratic state (i.e., for showing that the state can be an improvement). A state's citizens have unconditioned control if it is not conditional on the will of the state or anyone else. Influence over the behaviour of an agent is not adequate for control if this influence is conditional on the good-will of either the influenced agent or some other agent.

In the ideal stateless society, because all agents (including group agents) are perfectly morally motivated, they will not interfere in the choices of other agents in ways not accepted by those agents. However, individual agents are still dominated in Pettit's terms by collective agents because their influence on the collective agents' interference in their choices is entirely conditional on the good-will of the collective agents: if the collective agents *were* to choose to interfere against the will of an individual, the individual would be unable to prevent it. Thus, because individuals in the ideal stateless society do not have *unconditioned* control over the interference of collective agents, there is domination. It is essential, then, for the introduction of the state to count as an improvement, that its capacity to interfere be controlled *unconditionally* by its citizens.

How does Pettit think citizens of a democratic state can come to have *unconditioned* control of its actions? What can make the state's following the direction imposed on it by the democratic decisions of its citizens non-optional? Popular influence, he answers, has to impose an equally acceptable direction on government 'independently of the willingness of government to go along and independently of the willingness of any other agency to have the government go along'.[156] If the state has no choice but to follow the direction imposed by the people, then they have the sort of counterfactually robust control required. It has to be the case that if the state wanted to do otherwise it *could not*, or that there would be heavy costs associated with doing otherwise. It is not sufficient for control that the controlled agent's compliance be made *improbable*.

---

[156] Ibid. 172

Pettit considers a case in which you give me the key to your alcohol cupboard asking me to prevent you from accessing it until a specified time, and in which 'I go along with your instructions, but merely with a view to giving you the pleasant illusion of control; suppose that I am about to exit our relationship and think of this as a parting, somewhat sardonic gift'.[157] This does not count as your having control, he says; although my compliance with your direction is probable (or even guaranteed) *I* remain in control. In the normal alcohol cupboard case where I do as you ask out of friendship, you do have control because 'the relationship matters to me, whether for intrinsic or instrumental reasons, and there are heavy costs associated with refusing to go along'.[158] Thus, though probability of compliance is not sufficient for control, my having control over your interference does not require that it be *impossible* for you not to comply with my direction. There being heavy costs associated with non-compliance suffices. How can citizens of a state get the sort of control where non-compliance with their democratic direction is not merely improbable, but ineligible (or sufficiently costly) for the state? There are many ways in which constitutional and institutional design can make state usurpation of control *improbable*. This, though, is not what we are looking for.

Pettit says that the people can have this sort of robust control when there is 'a disposition of people to rise up in the face of a government abuse of legitimacy and a disposition of government to back down in response to the fact or prospect of such opposition'.[159] In other words, the possibility (or likelihood) of resistance, rebellion or revolution can allow the state's superior capacity for interference to be non-dominating. However, it seems a mistake to think that the sort of control that people can have over state interference by virtue of the threat of resistance allows the state to be non-dominating in a way that corporate agents in the ideal stateless society are not.

Firstly, the kind of resistance that could make following the people's direction non-optional is usually extremely difficult and costly. If you have the capacity to force the choice on me between accepting your interference and paying some very high cost, this seems just as problematic as the ordinary case of domination. Secondly, though, and more importantly, the sort of control that citizens of a democratic state have in virtue of the citizen body's tendency to resist does not seem to be adequate. It does not seem to give citizens any control over the state

---

[157] Ibid. 171
[158] Ibid. 171
[159] Ibid. 173

that they do not have over group agents in the ideal stateless society. The state has a capacity to interfere in the choices of each individual subject that is uncontrolled by that individual. The individual can (in a democratic state) exercise some influence over the direction taken by the state, but that influence is not unconditioned. The state has the capacity to ignore the individual's decision if it wants to. There is nothing the individual can do to ensure that the state has *no choice* but to take account of her influence, except in concert with a large number of others. If the state were to decide to persecute, disenfranchise or interfere arbitrarily with one or two individuals, there is nothing those individuals could do to prevent it except with the aid of a large number of their fellow citizens. If agent A has the capacity to interfere in the choices of another agent B and B has no control over that capacity so long as a third agent C (or group of agents C, D, E…) are willing to allow A to interfere with B, B seems to be just as dominated (or lacking in status freedom) as in Pettit's ordinary case of domination. Certainly, B does not meet the eyeball test with respect to A.

Whether or not we want to say B is dominated in such a case, the state in a resistive society is not any more controlled by its citizens than group agents are in an ideal stateless society. Group agents in the ideal stateless society dominate individuals in that society because, in virtue of their size, they have the power to interfere in the choices of individuals that those individuals cannot control (and the individuals do not have an equivalent uncontrolled capacity to interfere in the choices of group agents). The same is true of the state. As noted above, no individual has the capacity to remove the state's option of ignoring their influence. The individual *can* control the state's capacity to interfere in concert with others. But the same is true of individuals in the ideal stateless society. If others are willing to support an individual in resisting the interference of some group agent, they will together be able to prevent it from doing so. No group agents in the ideal stateless society have the uncontrolled capacity to interfere with the entire society (or large subgroups of the society). It is true that no group agent will in fact abuse their power, and we can assume the same of the ideal state that is introduced to the ideal stateless society. But these group agents in the ideal stateless society are counterfactually controlled by the society as a whole in just the same way that the ideal state is counterfactually controlled by its resistive citizen body.[160]

[160] Of course, states depend on their citizen body to function (including to have the capacity to interfere that they do); it might seem that group agents in the ideal stateless society need not. Citizens can resist their state just by withdrawing support. But the sense

4. Natural inequality

In the previous section, we allowed in group agents to the ideal stateless society. I argued that, although this might introduce status unfreedom, the introduction of a state to the ideal stateless society can do nothing to eradicate this. When we allow in natural inequalities to the ideal stateless society as well, the case is exactly the same. Natural inequalities similarly create *asymmetrical* domination. If A is naturally more powerful than B, then B will be dependent on A's will in a way unmatched by A's dependence on B's will. Again, just as the state can eliminate domination by group agents by controlling the interference of these group agents, the state can also eliminate the domination of powerful individuals by controlling their interference with others. For exactly the same reasons as above, however, the introduction of the state will not improve matters in terms of domination. Subjects of the state cannot have unconditioned control of *its* interference that is any better than the control they can have of the interference of powerful individuals (or group agents) in the ideal stateless society. Subjects of the state can control its interference when acting in concert with others, but in just the same way naturally less powerful members of the ideal stateless society can control the interference of more powerful members when acting in concert with others.

Assuming rough natural equality, then, we showed that it is *possible* for there to be a stateless society with no status unfreedom, and so that freedom *as such* does not require state enforcement a priori. We can now see that it is also not the case that freedom in societies involving group agents or natural inequality requires state enforcement a priori: the introduction of the state is no improvement on the ideal stateless society even when it contains group agents and/or natural inequalities.[161]

---

in which their support is relevant to the issue at hand is their allowing the state's interference with others to happen, *either* by actively participating in it, or simply by not preventing it. A sufficient number of people in the ideal stateless society can *prevent* any given group agent from interfering with other individuals, so in just the same sense, group agents in the ideal stateless society depend on the (at least tacit) support of other members of that society.

[161] It is also worth noting that once we introduce group agents or natural inequality, Simpson's arguments may apply. Even if the citizens of a state *could* have adequate control of it, domination *among* the citizens will no longer be fully symmetrical. Thus, if they did adequately control the state, the state's domination would be removed, but there would be domination among the citizens in its place: group agents or uncoordinated groups will have the uncontrolled capacity to interfere with individuals, and this need not be symmetrical since individuals will not be members of exactly equivalent groups or will not be of equivalent natural strength.

5. The state as omnilateral lawgiver

The preceding argument assumed that the sort of freedom that we should be concerned to protect is the sort of status freedom (and so, status equality) in which Pettit is interested. One reason for making the latter assumption was that, as Simpson shows, freedom as non-domination cannot be realised; there will always be domination or dependence on the will of others. However, a theorist like Pallikkathayil, for whom it is not non-domination that matters but just independence of the unilateral will of others, may reject Simpson's conclusions (or deny their relevance to freedom as independence), on the grounds that with the introduction of the state, the only remaining dependence is on an *omnilateral* legislative will. It is thus worth seeing whether the a priori argument can be made if freedom is *not* understood as status freedom, but in this latter way.

The Kantian notion of freedom on which Pallikkathayil's argument is based is similar to Pettit's, but not identical. Kantian freedom is independence from the will of others. To be free in this sense is to have the capacity to set and pursue one's own ends, which is to say, not to be subject to the ends of others. Even in an ideal stateless society in which determinate property rights are somehow established without a state, Pallikkathayil argues, people will be dependent on the unilateral wills of others. My capacity to set and pursue the ends I choose will be dependent on the unilateral decisions of others not to interfere. This is true even if the capacity of others to interfere with me is equivalent to my capacity to interfere with them. Because others *could* violate my property rights, these rights are dependent on their unilateral wills. And in just the same way, the property rights of others are dependent on *my* unilateral will.

The state can remove the *capacity* of others to interfere, or at least make interference highly costly for others, so that it is no longer the case that they *could easily* (and at little cost) interfere *if* they wanted to. We have seen, though, that domination, if understood simply as the uncontrolled power of interference, is unavoidable. It seems plausible that this will be the same for dependence on the will of others. If the state removes the capacity of its subjects to interfere with each other at will, they will become dependent on *its* will. I argued that subjects of a state cannot control its interference in a way that improves on the ideal stateless society, and Simpson argues that even if they can, it will be at the cost of becoming dependent on each other's will again.

However, Pallikkathayil will argue that subjects of the state need not be dependent on its will in a troubling way since it simply enforces an *omnilateral*

will. The assurance the state gives of others' respect of rights *does not reflect the unilateral will of any agent.*[162] In other words, even if dependence on the will of others is unavoidable, dependence on the *unilateral* will of others *is* avoidable. The assurance that state enforcement gives is dictated by an impartial decision procedure, one that reflects the will of all. However, if having assurance dictated by an impartial decision procedure that reflects the will of all is sufficient to alleviate dependence on the unilateral will of those who provide the assurance, then *state enforcement* is not necessary for independence. The necessary assurance can be achieved by members of an ideal stateless society following an impartial or omnilateral decision procedure: there is no a priori reason to suppose that this is impossible without state enforcement. If such a decision procedure is required for freedom (and if the requirements of universal freedom are central to the requirements of right, as Kant and Pallikkathayil certainly think they are), then perfectly morally motivated inhabitants of the ideal stateless society will willingly adhere to a freedom-respecting omnilateral decision procedure, if one is set up. There will then be perfectly good predictive assurance of rights respect (just as good as a state could provide) that is dictated by an omnilateral decision procedure. I am not aware of any a priori argument that it is impossible for such a decision procedure to be set up without enforcement.

Of course, even with a non-coercive omnilateral decision procedure, people would still make unilateral decisions about whether to observe and follow its decisions. Being perfectly morally motivated, they would always decide to follow these decisions, but their doing so would not be enforced, and so it would be *up to them.* I think it is for this reason that Pallikkathayil thinks assurance in an ideal stateless society must depend on the unilateral will of others. However, there seems to be no way of getting assurance without some dependence on the unilateral will of others. When we introduce a state with enforcement power as well as an omnilateral decision procedure, the assurance we have that others will observe the decisions of this procedure still depends on the unilateral decisions of either the state or the individuals who make up the enforcement agency. There are rules (made by the omnilateral decision procedure) instructing the enforcement officials as to what they can enforce on other people, so in this sense what they do is not up to them. But there can be rules like this in the ideal stateless society too. It *is* up to the members of the enforcement agency whether to follow these rules (the decisions of the omnilateral decision procedure), just as

---

[162] Pallikkathayil (2017) 38

it is up to the members of the ideal stateless society whether or not to observe its omnilateral rules.

One possible rationale for claiming that the state removes problematic unfreedom is the idea that it is not up to the individual members of the enforcement agency whether to observe the rules, since if one member disobeys, there will be other members who can enforce their compliance (or punish their non-compliance). Therefore, no individual can make a unilateral decision about whether to obey or not. But it will always be the case that some group of the members of the enforcement agency *could*, if they joined together, disobey the rules of the omnilateral decision procedure. So long as we admit the possibility of *polyadic* domination or dependence, as, following Simpson, I think we should, this constitutes problematic dependence. Polyadic dependence is of course not dependence on anyone's *unilateral* will, but, as Simpson powerfully argues, there is no plausible reason to think that dependence on the *multilateral* wills of several agents is necessarily less problematic from the point of view of freedom. It is, of course, very unlikely that members of the enforcement agency would exercise their collective power to disobey the rules, since there are coordination problems making it difficult. But it is also very unlikely that members of the ideal stateless society would disobey its omnilateral decision procedure since they are perfectly morally motivated. In addition, if the state itself is an agent then, even if all it in fact does is implement (enforce) the decisions of the omnilateral lawgiver, it will always have the *capacity* to interfere at will, so there will be *dyadic* dependence or domination as well. And, as I argued above, it is not possible for citizens of a state to achieve the sort of control that could alleviate dependence on its unilateral will.

6. Conclusion

In this and the previous chapter, I have argued that two prominent sorts of a priori argument for the legitimacy of a certain kind of state fail to establish their conclusion. Ripstein's Kantian assurance argument fails because it gives us no reason to think that an exclusive enforcer is necessary for assurance. Counterfactual dependence arguments, of the sort suggested by Pettit and Pallikkathayil, fail too because the introduction of a state cannot eliminate dependence on the will of others. While in an ideal stateless society we are dependent on the wills of other members of society or corporate groups, in a state society we are dependent on the wills of enforcement officials and of the state itself. Given this, there is little reason to think that anything about the idea of freedom itself demands state enforcement, or that there is any a priori argument

for the legitimacy of states that meet certain conditions. If state enforcement is necessary for freedom this is because of certain empirical and contingent facts about the world we live in. I have not demonstrated this, since all I have done is show that certain arguments to the contrary are not successful. However, I think that in the absence of a compelling a priori argument for state legitimacy, we can at least provisionally conclude that empirical considerations about the world we happen to live in are needed to establish the legitimacy of existing states.

Arguments like those discussed in the past two chapters *could*, perhaps, establish their conclusions if feasibility premises are added: premises such as the claim that the only *feasible* way of achieving status freedom or freedom as independence is through the state's exclusive enforcement. As I argued about the assurance argument, this premise will need to claim that the only feasible way of achieving freedom involves *this particular state's* exclusive enforcement, not just some state's exclusive enforcement; the same will be true for the counterfactual argument discussed in this chapter. In chapter 2 I argued that all of the arguments that might plausibly ground the legitimacy of some actual states depend on feasibility premises. Again, I did not demonstrate conclusively that any argument for the legitimacy of an existing state must depend on feasibility premises, but having surveyed prominent and plausible such arguments, it seems reasonable to conclude that we should not take the legitimacy of existing states for granted unless we can also take the necessary feasibility premises for granted. Thus, in order to get a picture of how feasibility premises can and cannot be used in moral argument, I turn, in the next chapter, to the question of how we should understand the concept of feasibility.

Chapter 5:
The concept of feasibility

The first part of this thesis claimed that plausible arguments for the legitimacy of existing states (their general and exclusive moral permission to enforce) depend for their success on premises about feasibility. They depend on premises like 'the only feasible way to achieve status freedom or freedom as independence is through the state in question's ('state S's') exclusive enforcement' or 'the only feasible way to achieve assurance that people will respect each other's rights is through state S's exclusive enforcement', or 'the best feasible world is one in which state S is an exclusive enforcer'. Whether premises like these can play the role that they need to play in these arguments will depend on when exactly infeasibility acts as a constraint on moral permissibility. Does something's being the only *feasible* way to do something morally required (or just morally good) mean that it too is morally required (or morally good)? Does something's being *infeasible* mean that it cannot be morally required, and thus that it is morally permissible not to do it? In order to answer these questions, and others like them, it seems that we will do well to have some idea of what exactly it means to say that some outcome is or is not feasible. How do we adjudicate these claims? What are their truth conditions?

It is quite common in moral and political philosophy to criticise and reject theories for making recommendations that are not feasible (call this a *feasibility critique*). A normative principle or recommendation cannot be true, or apply to us, it is commonly thought, if it is not feasible for us to do as it requires. If this is right, then it is quite plausible to think that, for example, if some action is the only feasible way to do something morally required, then it too is morally required. The question that we are interested in, then, is one that is of wider interest for moral and political philosophy. The account of feasibility I give in this chapter will be a general one, and so its conclusions will be relevant generally to the use of feasibility critiques in normative philosophy, but in the next chapter I will use it to draw some conclusions about state legitimacy in particular.

As noted above, the standard view of the relation of feasibility to moral theory takes feasibility to be a straightforward constraint on moral theory: proposals are

ruled out as morally required if not feasible.[163] For feasibility to operate as a constraint in such a straightforward manner (for there to be a single simple constraint on morality in general) there would need to be a single relevant binary notion of feasibility. In this chapter, I aim to argue that there is no *single* privileged binary concept of feasibility that can be assumed to play this role. Our ordinary concept of feasibility, I argue, is *multivocal*; that is, there are many different ways in which it can be made precise (which I call 'sharpenings'), with different truth conditions. Not *all* of these can plausibly be thought to constrain morality, and there is no immediate reason to suppose that any particular one of them is *the* constraint for morality in general. The import of my account of feasibility, then, will be to create a burden of proof. The thought that there is a simple and general constraint on moral theory cannot be justified by appeal to the ordinary concept of feasibility: we should not assume without argument that any *one* of the possible sharpenings of 'feasibility' is always a constraint on moral requirement. If there is a simple and universal constraint of this sort, it will need to be established by substantive argument (and if we cannot assume there is such a simple feasibility constraint on moral requirement, then something's being the only feasible way of doing something morally required and so being morally permissible is also no longer such a straightforward matter). Some thought, at least, needs to be given to *which* sharpenings of 'feasibility' constrain moral requirement *when*. Thus, to conclude that something is morally permissible because the only feasible way of φ-ing, we will need some reason to think that the sharpening according to which alternatives are *not* feasible is the appropriate one.

One way to reject the idea that the concept of feasibility is straightforwardly requirement-blocking would be to claim that no unified account of feasibility can be given at all. I think, though, that it *is* possible to give a unified account (that captures ordinary use), and I will argue for this simply by giving one. However, the best unified account going, I argue, is not *univocal* and so does not make feasibility claims *determinate*. There is a whole range of different 'sharpenings' of 'feasibility', a whole range of different binary definitions, but all of these can be unified under a single schema (or set of schemas). In the first part of the chapter, then, I will propose such a set of schemas in order to flesh out the multivocal account. Though there is no *single* binary definition of 'feasibility' as such, I will give an account of binary feasibility *given* a choice of sharpening. The account I

---

[163] Gilabert and Lawford-Smith (2012) and Lawford-Smith (2013) argue that the role of feasibility is not only to rule out proposals but also to contribute to ranking them. I do not intend to deny that it may have this latter role.

give is a *possibility*-based account: feasibility is to be analysed, I argue, in terms of possibility. To claim that some outcome is *feasible* is to make a claim about something being *possible*, but it is not simply to claim that that outcome is possible. Though related, possibility and feasibility are distinct concepts. Feasibility, unlike possibility, has to do with agency. It is about what it is possible for agents to do intentionally, in a way that my account will make clear, not just what might possibly come about.

I will defend the multivocal account (in the second half of the chapter) by showing that a variety of attempts to give a *univocal* account (or an account that gives 'feasibility' a small number of determinate senses) encounter problems that do not trouble a multivocal account. This does not yet motivate *my* (possibility-based) account, but the account given in the first half of the chapter is offered as a plausible and natural way of fleshing out a multivocal account. After motivating the need to go multivocal, I argue that there is no obvious plausible way to provide a *probability*-based multivocal account instead. My account, then, I argue, best captures our intuitions about feasibility. It might, though, be thought that this just shows that it is some concept other than feasibility that is the relevant constraint on morality. I thus argue, finally, that the alternative accounts of feasibility, not only do not capture the ordinary concept of feasibility, but also do not plausibly identify a general constraint on moral requirement.

1. *Feasibility and desirability*

Before I proceed to give an account of the concept of feasibility, it is important to note that feasibility critiques are frequently not as simple as the above description suggested. Often when a proposal or principle is criticised for being infeasible we do not really mean that it is infeasible simpliciter. Often these critiques are mixed up with questions of desirability, that is, with evaluative or normative questions. When we say that, for example, communism is not feasible, we may well not mean that it really is infeasible *as such*, but rather that it is not feasible in a desirable way. The concept of feasibility, as I understand it, is not an evaluative or normative one. However, sometimes when we say that some proposal is not feasible, this is shorthand for the claim that it is not feasible in conjunction with certain other things that we take to be more desirable or with the observance of other principles, which we take to be weightier than the proposal in question.[164] A feasibility critique of this form, then, says something like 'given that we should do

---

[164] Alan Hamlin (2012) makes the point that it is often through the infeasibility of *combinations* of things that feasibility constraints bind.

x or realise (values *v*, proposals *p*, principles *q*), it is not feasible to do/bring about/realise *y* in a desirable way'. Thus, when we say, for example, that X-ing is the only feasible way of doing something morally required, we may really mean that it is the only desirable (or morally acceptable) feasible way of doing the morally required thing.

The fact that feasibility critiques are often mixed up with questions of desirability does not mean that questions of feasibility themselves are evaluative or normative. The feasibility question is separate from the desirability question.[165] Or rather, the two *can* be separated, though often we put them together.[166] However, in general, what tends to be most important to know is not just whether some proposal is feasible, but rather whether it is feasible in an all-things-considered desirable way. We need to ask whether it is feasible in conjunction with the realisation of those other principles or values that would make it all-things-considered desirable. Nevertheless, questions about feasibility are independent of questions about morality.[167]

2. *Background*

Though feasibility has long been in the background of political philosophy, it has not until recently been explored in any detail. In an early philosophical discussion of feasibility, Juha Räikkä described what he took to be the standard binary approach to feasibility.[168] Proposals are feasible, on this approach, if and only if they are not rendered impossible by certain strong constraints. Feasibility began to become more prominent with the flowering of debate over 'ideal theory' and 'realism' in political theory. The paradigm of 'ideal theory' was criticised for being

---

[165] Gilabert (2008, 415), (2009, 663-4) and Gilabert and Lawford-Smith (2012, 816-7) also make this point.

[166] Juha Räikkä (1998) argues, on the contrary, that a definition of feasibility itself should include 'the necessary moral costs of changeover' as a constraint on feasibility. However, I think this is a mistake. The question of whether some outcome is feasible *in conjunction with* acceptable changeover costs is a different feasibility question to whether the outcome is feasible tout court.

[167] Brennan and Sayre-McCord (2016) argue that normative facts do make a difference to what is feasible, since normative judgments affect what people are willing to do and so how they behave, and these normative judgments are sometimes explained by their truth, by their reflecting the normative facts. However, the influence that the normative *facts* have on feasibility is *wholly* mediated by people's normative judgments. If we have the information about people's normative *judgments*, information about the normative *facts* would tell us *nothing* additional relevant to feasibility.

[168] Räikkä (1998) 32. He references as an example Beitz (1979) 156. (There were some earlier discussions of feasibility in the social scientific literature, such as Majone (1975), but nothing I am aware of in philosophy.) Some attempts to give univocal, binary definitions of 'feasibility' appear in Cowen (2007); Buchanan (2004); Hawthorn (1991, 158); Räikkä (1998) and Jensen (2009), and attempts to give a handful of binary definitions in Miller (2008), Brighouse (2004) and Elster (1985).

useless for action-guidance in the real world and an assumption underlying many of these arguments was that the recommendations of ideal theory, because it abstracted away from many of the facts of the world, were *infeasible*.[169] Geoffrey Brennan and Philip Pettit explicitly cast roughly this issue as 'the feasibility issue'.[170] These authors did not generally address in much detail what exactly it meant for a proposal to be *feasible* or *infeasible*. However, the importance that they saw feasibility to have gave the impetus for the development of such an account.

In the past ten years some attempts have been made to give a serious account of the concept of feasibility and what David Wiens has called the 'conditional probability' account has emerged as a prime contender.[171] In defending an '"ought" implies "feasibility"' principle, Geoffrey Brennan and Nicholas Southwood considered how 'feasibility' should be understood.[172] They reject a 'logical or nomological possibility' understanding of feasibility because there are many things that are logically or nomologically possible that they think do not qualify as feasible, such as a medical ignoramus performing a neurological operation for which he lacks the relevant expertise, or the realisation of a communist ideal. These things, they think, are not feasible because, though possible, they are not probable. It is *possible* that by sheer luck the medical ignoramus could perform the exact correct sequence of movements to perform a neurological operation. However, this is extremely unlikely, and we do not want to say that it is *feasible*. On the other hand, they reject the view that 'feasibility' just means 'sufficient probability'. If someone is too lazy to get out of bed at the weekend to go and watch their daughter's hockey games, then it may be very improbable that they will do so. However, it seems clearly not to be the case that this is thereby infeasible. On the contrary, it is perfectly feasible, but not likely. Thus, Brennan and Southwood opt for a conditional probability account of feasibility. The reason it is improbable that the lazy parent will go to their daughter's hockey match is that it is improbable that they will *try*. However, the conditional probability of their going *if they tried* is presumably much better. Thus, they say, feasibility should be understood 'in terms of reasonable probability of success conditional on trying'.[173]

Some of the most extensive treatments of the concept of feasibility, in particular those by Holly Lawford-Smith, Pablo Gilabert and Zofia Stemplowska, have

---

[169] E.g. Phillips (1985); Goodin (1995); Sen (2006) and (2009)
[170] Brennan and Pettit (2005)
[171] Wiens (2015a) 449.
[172] Brennan and Southwood (2007)
[173] Brennan and Southwood (2007) 9-10

followed along similar lines.[174] They, however, modify the account in ways that will appear below. David Wiens has, on the other hand, rejected the conditional probability account and defended what he calls a 'restricted possibility account', which is closer to the account I will give below in that it analyses feasibility in terms of *possibility* rather than *probability*.[175] In the next section I will set out a possibility-based, multivocal account of feasibility, and motivate it intuitively. I will then return to the rival accounts just mentioned and argue that they encounter problems to which my account is not vulnerable.

3. The multivocal account

One might assert 'It is not feasible to institute a system of participatory democracy'. Supposing that we know exactly what it would take for a system of participatory democracy to be instituted, there are still many different things this statement could mean; it is not clear, without any context, what exactly its truth conditions are. This, anyway, is what I will argue. In the second half of this chapter (section 4), I will argue that available *univocal* accounts of the concept of feasibility fail to capture intuitive judgments about feasibility, in a way that can be resolved by abandoning the attempt to identify a *single* binary notion of feasibility. In this section, I will flesh out in detail the multivocal account and explain how the various 'sharpenings' of 'feasibility' can be unified under a general set of schemas. The detailed account I give here is offered as a plausible and natural way to flesh out the multivocal account. I will motivate the various elements of my detailed account, but there may be alternatives, and I do not as such defend *this* account, except by arguing (in section 5) that there is no obvious way of giving a multivocal account instead in terms of *conditional probability.*

I proceed in this way, by first setting out the details of my multivocal account and leaving the argument for its multivocality until later so as to have a full fleshed-out account to compare to the various available univocal accounts. First, though, let me briefly intuitively motivate the idea that our ordinary concept of feasibility is a multivocal one. Return to the assertion mentioned in the first sentence of this section. I submit that there are various natural interpretations of this claim, with different truth conditions. Different ones may be salient in different contexts. Put roughly, one thing we might mean is that with the state being as it currently is, with people's motivations and preferences being as they currently are, and with the political system and power balance being as they are, we cannot institute

[174] Gilabert and Lawford-Smith (2012); Lawford-Smith (2013); Gilabert (2017) and Stemplowska (2016).
[175] Wiens (2015a)

participatory democracy. If this is what we mean, it seems likely that the statement is true. However, this seems not to be the only thing that might naturally be understood by this claim; in a different context, we might hold fixed a different range of facts. At the far extreme, we could mean that such a system is not physically, or logically, possible. If we meant *this*, we would probably say that the statement is false. This, perhaps, is an interpretation that will not be salient in many ordinary contexts. However, there seem to be various other interpretations in between, which might represent perfectly natural uses of the term 'feasibility'. We could mean, for instance, that, even if people's basic motivations, the power balance and so on are allowed to change, a system of participatory democracy is made impossible by some reasonably deep facts about human nature. There is no immediately obvious reason to suppose that any one of these various possible readings of the claim is privileged over the others as representing the 'proper' use of the term 'feasibility'. It would seem like a perfectly sensible response to the question 'is it feasible to institute a system of participatory democracy?' to ask, 'holding fixed what?'. (Of course, in many ordinary conversational contexts, some particular range of facts to hold fixed will be made salient, but if the question is asked with no relevant context, it is not obvious how to answer.)

I think that what this points to is the availability of a variety of different possible sharpenings of the term 'feasible', i.e., different ways in which we might make it precise. There is no obvious reason to think that any one of these is privileged as *the* thing we must mean by 'feasibility' *simpliciter*. Each sharpening, or precisification, corresponds to what I will call a 'feasibility constraint' (FC). An FC is a selection of which facts of the actual world to hold fixed (and correspondingly, which to allow to vary).[176] For each possible FC, there is a possible sharpening of 'feasibility'. Feasibility, then, is assessed (and feasibility claims have truth conditions) relative to a choice of FC. Often it is obvious from conversational context what FC, or rough range of FCs, is assumed in talking about feasibility. That is, sometimes in a particular context, a choice of which facts of the world to hold fixed is tacitly assumed and understood when a feasibility claim is made. However, this need not always be the case. Sometimes, when we make a feasibility claim we fail to say anything determinate because no choice of

---

[176] These can be thought of as what Kratzer (2012) calls 'conversational backgrounds', where conversational backgrounds are functions mapping possible worlds to sets of propositions (those propositions relative to which modal claims are to be evaluated for each world) (32-3). My account will thus bear some similarities to a Kratzerian account of modality, where the truth of modal sentences is relative to sets of propositions, but the only propositions that are relevant to *feasibility* are propositions true of the actual world (or the world for which feasibility is being assessed).

FC is understood or specified.[177] Some possible sets of facts (FCs) relative to which feasibility claims could be made may never be used in this way, such as perhaps the extreme cases where we hold fixed nothing but the laws of logic or where we hold fixed *everything*. It may seem like when we hold fixed so little or so much we are not really talking about *feasibility*, but I am happy to say that these are *limiting cases* of feasibility that may be little use for the sorts of purposes for which we ordinarily use the concept of feasibility. What seems clear is that there is at least a fair range of different sharpenings of the concept relative to which feasibility claims are and can be made.

I defend this picture below by arguing that attempts to give a *univocal* account of feasibility are subject to problems that can be resolved by going *multivocal* in the way just sketched. First, though, in this section I will fill out my multivocal account. Once a choice of FC has been made, then, what does it mean to say that some outcome is or is not feasible? There is no single binary concept of feasibility *tout court*, according to my multivocal account, but I think we can give a binary definition of feasibility *given* a choice of FC. If we simply claim that an outcome is feasible or infeasible, we fail to say anything determinate, but if a choice of FC is tacitly assumed or explicitly specified, any given outcome will be determinately either feasible or infeasible.

I think that the details of the multivocal account I am about to propose give a plausible and natural way of understanding feasibility claims. I start from the intuitive idea that feasibility is a special form of possibility and then motivate certain modifications in order to deal with problem cases. In the next section, I argue that an account of feasibility ought to be *multivocal*, and this supports the account I am about to give as the most natural candidate multivocal account. I suggest in section 5 that there is no obvious way to give a multivocal account in terms of *probability* instead.

3.a) Feasibility on an FC

In order to give such a binary definition of feasibility *given* a choice of sharpening, the first question to ask is what feasibility is *of*. It seems clear that feasibility can be of outcomes, or states of affairs. We might think, though, that actions can also be assessed for feasibility. We might wonder whether it is feasible, say, for me to

---

[177] We could think about feasibility in terms of a supervaluational structure, meaning that if a proposal is 'superfeasible' (feasible on all FCs) then we can say straightforwardly that it is feasible *tout court* and if it is 'superinfeasible' (infeasible on all FCs) then it is infeasible *tout court*. Thus, the question may have a determinate answer if the proposal is either superfeasible or superinfeasible, but most of the time this will not be the case.

run to Africa. I think actions can certainly be assessed for feasibility, but actions can *be* outcomes. That is, for every action φ there is an outcome consisting in X's performance of φ. Thus, for the sake of simplicity, we can bring all the categories that can be assessed for feasibility under the category of outcomes. The left-hand-side of my definition, then, will be the schema, modified from Gilabert and Lawford-Smith:

O is feasible (for X) in Z on *f*.[178]

where O is an outcome, Z is a context (the set of facts of the world that we start from in defining an FC – in other words, a choice of time and possible world; FCs are defined as subsets of the facts that hold at this time in this world, facts which are to be 'held fixed') and *f* is some FC.

Now, G.A. Cohen has suggested that there are two elements to feasibility: accessibility and stability.[179] Accessibility is a matter of whether we can get there from here, while stability is a matter of whether the outcome can be maintained if we do get there. Sometimes 'feasibility' is used simply to mean accessibility but in other uses it requires *both* accessibility *and* stability.[180] (Really what is required is that the outcome be *capable* of being stable.) These two things are separate: an outcome can be accessible but not capable of being stable, or capable of being stable but not accessible. I will focus for now on accessibility, but all that will be needed to get an account of the use that requires stability is to add to the account given below a requirement that the outcome be brought about *stably*, and an account of what stability is, which I will give below.

3.b) Agent-Relative Accessibility

I think that, most plausibly, accessibility given a choice of sharpening is a matter of possibility restricted in various ways. This will be motivated only intuitively, as well as by the lack of any obvious way of making a multivocal account work in terms of probability. If it is asserted that some outcome is accessible for us, it is

---

[178] Gilabert and Lawford-Smith (2012) 812.
[179] Cohen (2009) 56-7. This distinction is very similar to Erik Olin Wright's (2006) distinction between 'viability' and 'achievability' (97-9).
[180] For this reason David Wiens (2015a) argues that stability is not a necessary condition for feasibility (3, n. 2). It seems clear that there is a use of 'feasibility' for which stability *is* a necessary condition, as in when we say that communism is not feasible because human nature will lead it to collapse. Gilabert and Lawford-Smith argue that getting to some outcome, if it cannot be maintained, does not really look like 'getting there' at all (2012, 813). Wiens is right, though, that this is not always how 'feasibility' is used, as in when we say that it is feasible to balance a spinning-top on its point, despite the instability of that position.

natural to think about the truth of this as having to do with whether it is *possible* for us to get there, or to bring it about. What is accessible for us, goes the thought, though not simply anything that might possibly come about, has to do with what might possibly be brought about by us. Brennan and Southwood begin their discussion of feasibility by rejecting the two simple and natural suggested analyses as simple possibility or simple probability.[181] I start with possibility, since, as I argue below, there is no straightforward way to give a *probability*-based multivocal account (and because it seems to me more intuitively plausible to associate accessibility with possibility than probability). Brennan and Southwood reject a simple possibility account on the basis of counterexamples, but, as we will see, a more sophisticated possibility account is not subject to such counterexamples.

On the possibility-based account I will present, then, feasibility is not simply equivalent to possibility: something's simply being *possible* is not sufficient for its being feasible. But feasibility can be cashed out *in terms* of possibility.[182] For something to be feasible, as already noted, it needs to be possible for it to come about in a particular way, one that involves agency. I think this is the most intuitively plausible way of understanding what feasibility is: it is about what we *can* bring about, what it is *possible* for *us* to *bring* about. For something to be feasible given a set of facts being held fixed is for it to be possible (in the restricted way) compatibly with these facts.

Thus, I propose the following definition for binary accessibility given a choice of FC:

> *Agent-Relative Accessibility.* O is accessible for X in Z on f if and only if $\exists\varphi$(X's φ-ing to bring about O in Z is possible given constraint f, that is, is not incompatible with constraint f)

> where 'φ-ing to bring about O' means performing some intentional action φ that will bring about O, or will make things such that an event e occurs that will bring about O, (though it need not be intended to bring about, or contribute to bringing about, O) such that X brings about O *safely* and *competently* (notions to be explained below).

In order to see what is involved in something being possible given some FC, it may help to think of an FC as playing a similar role to an accessibility relation in

---

[181] Brennan and Southwood (2007) 8-9
[182] Gilabert and Lawford-Smith (2012) define binary feasibility in terms of possibility compatibly with certain constraints: I think we just need to extend this to *any* set of constraints, rather than merely the hard constraints they identify (as I will argue below).

modal logic.[183] An event is possible given an FC if it occurs in some possible world out of a restricted range selected by the choice of FC and Z.[184] The world from which the accessible worlds must be accessible (call this the home world) is selected by Z (it is likely to be the actual world, but need not be). The accessible worlds are then restricted to those identical to the home world up until the time of Z. Finally, the FC then restricts the accessible worlds to those in which, after that time, all facts remain fixed except for those that the FC allows to vary. If an outcome is brought about in the right way (directly or indirectly) by X in some possible world out of this restricted range, then it is accessible for X given this FC (and Z). What this means, in less abstract terms, is that when we choose a range of facts to hold fixed, say the deepest facts of human nature along with the laws of physics, biology and so on, an outcome is accessible for me if and only if there is some possible world in which those laws and facts of human nature hold (and which is identical to the actual world up to now) in which I bring about the outcome in question.

Although, as I have said, feasibility and possibility are distinct concepts, and what is distinctive about feasibility is that it has to do with what *agents* can bring about through intentional action, my definition does not require that for O to be accessible to X it must be possible for X to bring about O *intentionally*. It is possible, I think, for certain outcomes to be feasible for an agent despite its not being possible for the agent to bring them about *intentionally*. This appears to be brought out by the case of a university student, Florence, taking an exam, who is unfamiliar with the grading system for that exam.[185] The student's teachers might ask whether it is feasible for Florence to get a 2:1 in the exam. It seems like it *can* be feasible for her to get a 2:1 even if she has no idea what a 2:1 is and is simply intending to do as well as she can on the exam and find out later what mark this translates into. If she does not know what a 2:1 is, she cannot intentionally bring about her getting a 2:1, but she can intentionally act in a way that results in her getting a 2:1 and it seems like in some cases this can be sufficient for an outcome's being feasible for an agent.

---

[183] Wiens (2015a) also suggests thinking of feasibility in terms of an accessibility relation on possible worlds (457), but his account fails to deal with all of the problems that motivate my account (see below).

[184] Note that *e* must be an event that occurs in one of these possible worlds and brings about O. It may be synchronically or only diachronically possible (on this distinction see Jensen, 2009). That is, it must just occur at some point in one of those possible worlds, it need not be immediate from the time of Z.

[185] Thanks to Han van Wietmarschen for suggesting this case to me.

It is not enough, though, that it be possible for X's intentional action to bring about O for O to be feasible for X. We have already seen that such a simple possibility account will not do. It is *possible* that a medical ignoramus who sets out to perform brain surgery by trying out a random series of movements will choose exactly the right sequence of movements and so successfully perform brain surgery. (This is, of course, possible even consistently with the medical ignoramus's knowledge, skills and so on being held fixed.) It is surely not, though, thereby *feasible* for the medical ignoramus to perform brain surgery. It needs to be possible for the agent to bring about the outcome *not just by freak luck*. For this reason, I add the requirements of *safety* and *competence*. I borrow these notions from the literature in epistemology, where it is often thought that an account of knowledge must accommodate the intuition that true belief achieved by luck does not count as knowledge.[186]

Firstly, O's being feasible for X requires that there be a possible action of X's that brings about O *competently*, by which I mean that O is creditable to some sufficient extent to X's relevant *competence*. Ernest Sosa describes a competence as 'a disposition, one with a basis resident in the competent agent, one that would in appropriately normal conditions ensure (or make highly likely) the success of any relevant performance issued by it'.[187] I will not attempt to give a full account of what a competence is, but will assume that there is an intuitive notion pointed at by Sosa's description. A competence is like a skill: some actions that an agent performs manifest competence or skill, while others do not. An experienced archer hitting a target in ordinary conditions and with no intervening factors seems to be an example of the former, while a game-player rolling a 6 seems to be an example of the latter. It seems clear that a medical ignoramus is *not* competent to perform brain surgery, though they *may* be competent to perform the precise sequence of movements that would be needed in a particular instance to perform brain surgery. (For this reason, the requirement is that the agent be possibly competent to *bring about the outcome*, not to perform the action that brings about the outcome; the medical ignoramus is competent to perform a sequence of actions that would together bring about the outcome of her performing brain surgery, but she is not competent to perform brain surgery.) On the other hand, Florence, the student in the above exam case, *may* be competent to get a 2:1 despite not knowing what this means.

---

186 See for instance Ichikawa and Steup (2017), Sosa (2007) and Pritchard (2012).
187 Sosa (2007) 29

My account requires that there be a possible world in which X's bringing about O is *sufficiently creditable* to X's competence. Sosa's account of knowledge demands that the correctness of a belief be *explained* by an agent's competence in order to count as knowledge. Duncan Pritchard argues that the requirement on knowledge is a weaker one: that the agent's true belief be the *product* of her relevant competence (but he also adds an additional safety requirement).[188] He motivates this with a case that he thinks Sosa's strong requirement will get wrong:

> *Jenny.* Jenny gets off the train in an unfamiliar city and asks the first person that she meets for directions. The person that she asks is indeed knowledgeable about the area, and helpfully gives her directions. Jenny believes what she is told and goes on her way to her intended destination.[189]

In this case, Pritchard thinks, Jenny's knowledge is not *explained* by any competence, or cognitive ability, of hers. It is, though, he says the *product* of her cognitive ability. The difference is that 'Jenny's cognitive success is not *primarily* creditable to her cognitive agency', while 'her safe true belief is *to a significant degree* creditable to her cognitive agency'.[190] A good account of feasibility should also be able to deal with cases like *Jenny*; we will want to say that it is feasible for Jenny to find out how to get to her destination (even holding fixed her knowledge, that of those in the area, and so on). Thus, we should not require that there be a possible world in which O is *primarily explained* by X's competence. But how far an outcome's being brought about is creditable to an agent's competence is a matter of degree: there must be some degree of (possible) creditability to X's competence that is sufficient to make an outcome feasible for X. I will not attempt to determine exactly what this degree is. (Note that it will not be sufficient merely for X's competence to have *some* (possible) role in bringing about the outcome. Imagine that X is a competent dart thrower who throws a dart headed towards the bullseye. However, a malevolent onlooker, Y, is poised to blow X's dart away from the bullseye as it approaches. Fortunately, though, there is another, benevolent, onlooker, Z, ready with a wind-machine to direct the dart back to the bullseye. X throws the dart, and, after the interventions of Y and Z, it hits the bullseye. X's competence plays *some* role in explaining the dart's hitting the bullseye, but if this is the only possible way in which X can hit the bullseye, it is not plausibly feasible for her to do so.)

The inclusion of a requirement of (possible) competence does not turn my account into a dispositional account of feasibility (one that analyses the feasibility of O for

---

[188] Pritchard (2012) 273
[189] Ibid. 269, adapted from Lackey (2007).
[190] Ibid. 273-4, emphasis added.

X as X's being *disposed* to bring about O in certain conditions). The requirement of my account is not that X *actually* be competent to bring about O, but rather that it be *possible* for X to be so competent. There must be a possible world in which X performs an action that brings about O, and in which O is sufficiently creditable to X's competence. This, of course, requires that X *be* competent to bring about O, but this requirement applies to the possible world in question. No doubt to say that X is competent to bring about some outcome in a world *w* is to make some claim also about worlds other than *w*. For instance, it may entail that in a certain proportion of the worlds close to *w*, those in which circumstances are relevantly similar, X brings about O. However, this need not entail anything about the *actual* world.[191]

I also add a safety requirement. The requirement is that it must be possible for X to bring about O safely, where the notion of safety is again borrowed from the epistemology literature.[192] Sosa characterises safety thus: 'A performance is safe if and only if not easily would it … have failed, not easily would it have fallen short of its aim'.[193] For a performance to be safe, it needs to be the case that it succeeds not only in the actual world, but also in other nearby worlds, similar to the actual world in certain relevant respects. If there is a possibility sufficiently close to the actual world in which the performance does not succeed, then it too easily could have failed. The requirement on feasibility is that there be a *possible* world *w* in which X brings about O safely. This means that in all the sufficiently close possible worlds to *w*, in which circumstances are relevantly similar, X succeeds in bringing about O. The requirement is not that in all those close possible worlds *in which X φs*, X successfully brings about O, but just that in all those sufficiently close possible worlds in which circumstances are relevantly similar, X brings about O. The addition of a safety requirement on top of the competence requirement is needed because there could be cases where some piece of freak luck makes possible the exercise of a competence. We would not want to say that it is

---

[191] What exactly the disposition is that is involved in a competence for bringing about O, in what sense it must have its basis in the competent agent, and what exactly is required for an outcome to be creditable to an agent's competence are questions that I will have to leave unanswered. The disposition cannot simply be a disposition for the agent's φ-ing to bring about O, since there is a possible φ that the medical ignoramus can perform that will tend to result in brain surgery being performed (a specific sequence of movements). It would presumably have to be something like a disposition to produce O given similar circumstances and in response to similar stimuli (where these include something like the agent's desires or preferences); cf. Pritchard (2012) 256-7.
[192] Note that the requirement is that *the agent* bring about O safely, not that the action do so. For discussion of safety in epistemology, see for example Ichikawa and Steup (2017), Sosa (2007) and Pritchard (2012).
[193] Sosa (2007) 25

feasible for X to bring about O where the *only* way that X could possibly bring about O is where some piece of freak luck allows her to exercise a competence. For example, suppose there is a brick wall separating Ella, a competent darts player, from a dartboard. There are some large rocks in the vicinity that, if they were positioned in a particular spot, would allow Ella to climb on top and throw a dart over the wall at the board. However, the rocks are too heavy to move. The mere fact that it is *possible* that, say, a small landslide could happen to shift one of these rocks into exactly the right position, enabling Ella to exercise her dart-throwing competence, is surely not sufficient to make it feasible for her to hit the bullseye. This is because the conditions could too easily not have permitted Ella to exercise her competence. The possible successful bringing about of an outcome (hitting the bullseye) is too modally fragile, it is not *safe*.

There is some overlap in the work that can be done by the safety and competence requirements (they can deal with some of the same cases), but the safety requirement will not do on its own. The need for the competence requirement is brought out by a case of Sosa's:

> A protecting angel with a wind machine might ensure that [an] archer's shot would hit the bullseye … and a particular shot might hit the bullseye through a gust from the angel's machine, which compensates for a natural gust that initially diverts the arrow.[194]

The archer's shot hits the bullseye *safely* in this case, but the possibility of this scenario is not plausibly sufficient to make hitting the bullseye feasible, given the initial natural gust of wind. For this reason, the competence requirement is necessary as well.

3.c) Non-Agent-Relative Accessibility

Now, the above definition is a definition of *agent-relative* feasibility (or accessibility). This means that it defines the feasibility of an outcome *for some agent*. We may also, however, want a non-agent-relative definition of feasibility (on a given FC), a criterion for what it would take for an outcome to be feasible *tout court* on some FC (as opposed to feasible *for some X* on a given FC).[195]

I suggest the following:

---

[194] Sosa (2007) 29
[195] We can also talk about feasibility for some *group* of agents. My definition of non-agent-relative feasibility below gives an account of what it is for something to be feasible for the group of *all* agents. This definition can thus also give us a definition of feasibility for any particular group of agents, just by narrowing the domain (over which the quantifiers range) to the group in question.

Non-Agent-Relative Accessibility. An outcome O is accessible in Z on *f* iff
∃*X*∃Φ(X's Φ-ing is possible given constraint *f*, that is, is not incompatible with
constraint *f*).

where either
a) *X* is an agent and Φ is an intentional action that will bring about O (or will
make things such that an event *e* occurs that will bring about O) such that X
brings about O safely and competently, or
b) *X* is a group of agents and Φ is a set of intentional actions and for X to Φ is for
it to be the case that, for each *φ ∈ Φ* there is some *x ∈ X* such that *x* φs safely
and competently (see above[196]); and the combined result of all *φ ∈ Φ* is O; and
the group *X* brings about O safely.[197]

One might think that a non-agent-relative definition of accessibility ought not to
involve reference to agents and actions at all; an action is accessible on some FC
just if there is a possible event that would bring it about compatibly with that FC.
However, as I have said, I think that feasibility is a modal concept about *doing*. If
something is possible, but cannot be brought about by intentional action, then it
is not *feasible*, it is merely *possible*. This I take to be a key distinction between the
two concepts, feasibility and possibility; the former requires agency while the
latter does not. So, for an outcome to be (non-agent-relative) *feasible* (accessible) it
must be possible for it to come about in an *agential* (and intentional) way.

Again, this definition can be understood in terms of possible worlds. The
existential quantifier quantifies over the restricted set of possible worlds selected
by Z and the chosen FC together. The requirement of the definition is that there
be an action (or set of actions) that brings about O in at least one of these possible
worlds. To illustrate, then, participatory democracy is accessible on an FC that
holds fixed certain deep facts of human nature only if, compatibly with those facts,
it is possible for some agent(s) to bring it about in the right way. If there is an
agent or group of agents who, in any of the possible worlds in which the facts of
human nature hold, bring about participatory democracy (in the right way) then it
is accessible on this FC. What is necessary is that there is an agential route to the

---

[196] The definition of agent-relative accessibility above requires that the agent safely and
competently *bring about the outcome*. Here, non-agent relative accessibility requires rather
that each *x ∈ X* safely and competently *φ*. That is, each agent's performance of the action
φ must be creditable to a sufficient degree to their competence and would not too easily
have failed. It is then added that the *group, X*, must safely bring about the outcome O. The
considerations that motivated including a competence requirement in the agent-relative
definition motivate including one here also (at the individual level). Competence is not
required at the group level because it is not clear whether it makes sense to talk about a
group, which need not be an *agent*, or anything more than a simple *set* of individuals, being
competent to bring about a certain outcome. A *safety* requirement, though, is added at the
group level for reasons that will be discussed below.
[197] Cf. Lawford-Smith (2013) 247.

outcome; if the outcome is possible but only through non-actions then, though it is possible, it is not feasible.

Zofia Stemplowska raises a problem for accounts of feasibility: that of dealing with cases of joint, uncoordinated, action.[198] Some outcomes that require joint, uncoordinated, action seem to be infeasible even though it is *possible* for everyone to perform the action required, and even though if everyone were to try to bring about the outcome, they would likely succeed. In Stemplowska's example, if everyone were to try to touch their nose next Tuesday, they would likely succeed, but we might still think that the outcome consisting in everyone touching their nose next Tuesday is infeasible because of the difficulty of coordination. Stemplowska's response to this problem is to argue that what is important is that the agent *know how* to do the necessary action, where knowing how to do an action requires knowing that some action will bring about a given outcome.[199]

My account does not require that the agent *know how* to bring about the outcome in question, since it seems that an outcome *can* be feasible for an agent despite the agent *not* knowing how to bring it about. For instance, as in the case discussed above, it can be feasible for Florence to get a 2:1 despite not knowing what a 2:1 is.[200] Nevertheless, knowhow is not irrelevant to feasibility on my account. Though competence for bringing about an outcome O does not *require* knowledge how to bring about O, it will often, or usually, be the case that knowhow of some sort is bound up with competence. One can be competent to, say, get a 2:1 without knowing how to do so (because, perhaps, one does not know what a 2:1 is), but it seems clear that this competence will require knowledge how to do certain things. For instance, one will probably need knowledge how to write, how to sit an exam, answer the kind of questions asked by the exam, and knowledge of the material assessed. Thus, an agent's knowledge (both propositional and, especially, procedural) will, when held fixed by the chosen FC, constrain feasibility.

However, this does not deal with Stemplowska's problem, since for cases of feasibility for groups, the competence requirement only applies to the *individual* actions that make up the group action (or collectively produce the outcome). The problem with everyone in the world touching their noses next Tuesday is not that

---

[198] Stemplowska (2016) 276-7
[199] Ibid. 284
[200] My account does, though, require intentional action. Stemplowska (2016) seems nearly right when she says that 'the feasibility of actions depends on there being an intentional agent, single or collective, who can perform the action in question' (289). This is not quite correct, since in cases of joint, uncoordinated action there is no collective agent that performs the action, but the thought is close to the truth: there must be a set of intentional actions that bring about the outcome in question.

any individual agent is not competent to touch their nose. This problem is instead dealt with by the requirement that the *group* possibly bring about the outcome *safely*. The requirement here is that the group would not too easily have failed to produce the outcome. There has to be a possible world in which the group produces the outcome and in which it is true that in relevantly similar circumstances and in response to relevantly similar stimuli, the group would have produced the outcome (in all relevantly similar possible worlds, the outcome must also be produced by the group). It seems that, if we hold fixed the knowledge of everybody in the world, and their lack of coordination, there will be no such possibility in which the outcome in which everyone touches their nose next Tuesday is *safely* produced. There is, to be sure, a possible world, consistent with these constraints, in which everyone touches their nose next Tuesday. But it seems like the only possible worlds in which this happens will be ones in which it happens by *luck*, and thus not safely. Often, in fact, safety will require knowledge. If no members of a group possess knowledge how to bring about an outcome (or something equivalent to the outcome), it is likely that that group will not be able to bring about the outcome safely or reliably.[201]

3.d) Stability

I noted above that there is a use of 'feasibility' according to which an outcome must be both accessible *and* stable. To get a definition of this use, all that we need do is replace 'bring about O' in the above definitions of agent-relative and non-agent-relative accessibility with 'bring about O stably'. What, though, does it mean for an outcome to be *stable*? One writer who discusses this is Rawls, who put some importance in the stability of his conception of justice.[202] He defines stability for *systems*, whereas what I want is a definition of stability for outcomes, or states of affairs. However, what he says for systems will be useful as a point of departure. He says that stability for systems is a matter of the forces in the system that will return the system to *equilibrium*. A system is in equilibrium 'when it has reached a state that persists indefinitely over time so long as no external forces impinge upon it'. An equilibrium is stable 'whenever departures from it ... call into play forces within the system that tend to bring it back to this equilibrium state'.[203] Rawls thus requires that departures from a stable system must *themselves bring*

---

[201] Note, though, that my account does not, as Stemplowska's does, say that everyone's touching their nose next Tuesday is straightforwardly feasible. If we hold fixed agents' knowledge and lack of coordination and so on, I think, then it does. But on other FCs, in which these things are allowed to vary, it may in fact be feasible.
[202] Rawls (1999a) 398-400
[203] Ibid. 400

about a return to equilibrium. I don't see, however, why we must require this. Presumably an outcome would be stable whether departures from it tended to *bring about* returns to equilibrium, or whether departures were simply followed by returns.

Brennan and Pettit's discussion of feasibility centres exclusively around stability.[204] They suggest that institutions are stable if they can be kept in place and *enabled to promote the benefits for which they are designed.*[205] However, I want to reject this latter part, since stability, like feasibility, I take to be a non‑moral notion; this definition conflates the desirability and feasibility questions into one. Whether it will promote the benefits for which it is designed should be a different question. It is the question of whether the proposal *in conjunction* with the realisation of those benefits is stable.

Thus, my definition of stability is the following:

> *Stability.* An outcome O is stable in Z on *f* if, and only if, it can be sustained indefinitely as an equilibrium, compatibly with *f.*

The inclusion of the phrase 'as an equilibrium' is intended to indicate that a stable outcome need not be sustained indefinitely *and perfectly*, without any departures. There may be some departures from the given outcome, so long as they tend to be followed by a return to the outcome in question. There is obviously some vagueness here, since the question how frequent, extensive or pervasive departures from an outcome must be before we determine that that outcome is not stable is not given any clear answer.

I think that to be *stable* simpliciter an outcome must be sustainable *indefinitely*. However, I can see that we might want to say that outcomes that are not sustainable indefinitely, but for relatively long periods of time approximate more to stability. There is clearly a scale of unstable outcomes that approximate more or less to stability. An outcome that can only be sustained for, say, a day is further from being stable than one that can be maintained for long periods (years, decades perhaps), but that will eventually collapse.

Thus, this definition requires that for an outcome O to be stable on an FC *f*, there must be no facts held fixed by *f* that prevent O from being sustained indefinitely.

---

[204] Brennan and Pettit (2005)
[205] Wright (2006) also gives a moralised definition of his notion of viability, but it could be that his notion of viability is not simply equivalent to stability, but rather the combination of stability and desirability: it requires that outcomes 'actually generate in a sustained manner the emancipatory consequences that motivated their proposal' (97).

If some fact held fixed on *f* will tend to produce uncorrected departures from O, then O will not be stable. Thus, for example, on a certain FC that holds fixed a number of facts about human motivational nature, social life without coercion might be said to be unstable because if no-one coercively enforces law, then human nature (as held fixed on *f*) is such that people will be tempted to coerce others for their own advantage and so coercion will reappear.

4. Alternative accounts

The above concludes the presentation of my account of feasibility. It has two main features: first, there is no single binary concept of feasibility, but rather a range of possible binary sharpenings, no single one of which is obviously privileged over the others, and second, on any given sharpening, feasibility is a matter of what it is *possible* for agents to bring about safely and competently and through intentional action. I will now defend this account by showing that alternative accounts available in the literature are problematic. More specifically, I will defend the first of these two features, the 'multivocality' of my account. In this section, I argue that the main alternative accounts encounter difficulties that do not trouble a multivocal account of the sort offered above. This will also, however, indirectly constitute a defence of the second feature, the possibility-based account of feasibility *given* a choice of FC, since that account was offered as the most natural way of filling out a multivocal account, and avoiding the problems for univocal accounts that this section will adduce. In the next section, I argue that there is no obvious way to make a multivocal account work instead in terms of *probability.*

4.a) Conditional probability account

I will begin with the conditional probability account, which seems to be the most prominent account in the literature. This is the account suggested by Brennan and Southwood in response to the failure of the simple possibility and simple probability accounts.[206] It says that feasibility is a matter of reasonable probability of success conditional on trying. This analyses feasibility claims as counterfactual statements. The claim that it is feasible for A to bring about O becomes: 'if A tried to bring about O, A would probably (with a reasonable degree of probability) succeed'. However, given its simplest and most natural reading, this does not correspond to the ordinary concept of feasibility.

---

[206] Brennan and Southwood (2007)

We can begin to see what the problem is when we note that the proposed analysans is a straightforward counterfactual statement of the form 'If A were the case, C would be the case'. David Lewis's is the most well-known analysis of counterfactuals. According to him, a statement of this form is true at a world i 'iff some (accessible) AC-world [world at which both A and C are true] is closer to i than any A¬C-world [world at which both A and not-C are true], if there are any (accessible) A-worlds [worlds at which A is true]'.[207] Thus, the above statement could be analysed as:

> 'There is some (accessible) possible world at which A tries to bring about O and has a reasonable probability of success that is closer to the actual world than any world at which A tries to bring about O and has an insufficient probability of success'.

However, the closest possible world in which I try to, say, run a mile is not one in which I try very seriously. I am not very fit and am quite lazy, so a world in which I seriously try to run a mile (say, train myself, get fit, make repeated efforts) departs more from the actual world than one in which I just make a half-hearted attempt once and then give up. On the above analysis, then, it is straightforwardly *infeasible* for me to run a mile. But it is not obvious that it is: there seems to be at least a sense available in which it *is* feasible for me to run a mile, despite the fact that in the closest possible world in which I try, I fail.[208]

A plausible version of the conditional probability account, then, must demand probability conditional on something like *wholehearted* trying. This is obviously not a plausible account of the ordinary use of 'trying', but we could just stipulate that this is what is meant by the conditional probability account. My wholeheartedly trying to run a mile would no doubt involve training, making repeated efforts and so on, and so understood thus, the account would no doubt say that it is feasible for me to run a mile (at least within, say, a year). This gives us the following account:

---

[207] Lewis (1973) 424-5

[208] The account could alternatively be read as interpreting the feasibility claim as saying something like 'The probability that A brings about O given that A tries is sufficiently high'. If conditional probability is understood in terms of proportion of possible worlds, this reading could avoid the need to talk about probabilities in different possible worlds (which requires us to be able to make sense of single-case probability). We would simply take all of the possible worlds (or perhaps all of the sufficiently close possible worlds) in which A tries to bring about O and ask whether A succeeds in a sufficiently high proportion of them. However, there is an infinite number of possible worlds in which A tries. Thus, unless A succeeds in none or all of them, it is far from obvious what the proportion will be. In any case, if we *can* make sense of probability in this way, the same objections I make below to the counterfactual conditional probability account apply equally here. Southwood (2016) suggests moving from a counterfactual to a dispositional account of ability to avoid counterexamples to do with 'finkish dispositions': again, the objections I will raise apply *mutatis mutandis* to such an account of feasibility.

(CP) It is feasible for A to bring about O if, and only if, if A *wholeheartedly* tried (A tried her best) to bring about O, she would probably bring about O (or would *tend* to bring about O).[209]

What, then, is meant by 'wholeheartedly trying'? One option would be: 'performing the objectively best bundle of actions for bringing about O'. This, though, would give us a view quite different from the conditional probability account we started with, and one that loses some of the advantages that motivated it in the first place. It would say the following:

> (CPa) It is feasible for A to bring about O if, and only if, if A were to perform the objectively best bundle of actions for O (which presumably means something like: the bundle of actions, of those possible for A to perform, that would give O the highest probability), she would probably bring about O.[210,211]

This is open again to the kinds of objection Brennan and Southwood made against the simple possibility view. A medical ignoramus, if they performed the best bundle of actions (not just what they *consider* to be the best bundle of actions), would have a good probability of performing a neurological operation. The best bundle of actions is the sequence of manoeuvres that constitutes the neurological operation in question. The problem is that they do not *know* which actions to perform and so, in actual fact, if they tried, they would have a very low probability of success.

Alternatively, and perhaps more plausibly, we could understand wholehearted trying in terms of what the agent *believes* to be best. David Estlund has suggested to me that we could understand it as 'pursuing whatever is believed to be an effective (or likely to be effective) means to the outcome'. This would give us the following:

> (CPb) It is feasible for A to bring about O if, and only if, if A were to pursue whatever is believed to be an effective (or likely to be effective) means to O, she would probably bring about O.

It seems like the conditional probability account will need to be fleshed out in something like this way, and this is the strongest version of it that I am aware of. However, there are still intuitions that this account does not capture. The intuitions I will point to are cases where our intuitions pull in different directions, so the account's forcing us to go one way and not the other would not normally

---

[209]Stemplowska (2016) suggests understanding 'trying' in a similar way (275).
[210] Lawford-Smith's (2013) account of *scalar* feasibility is something like this: she makes the *degree* of feasibility of O equal to its probability given the best action (255).
[211] If probabilities are entirely subjective, then 'the objectively best bundle of actions' should be understood, rather, as something like the bundle of actions that a fully-informed, ideal observer would choose to bring about O.

count against it. However, in all of these cases, my account will accommodate the intuitive pull in both directions, where the conditional probability account fails to do so. Of course, the defender of the conditional probability account might accept these intuitive judgments about the use of 'feasible' but say that they are looking only for the concept that constrains moral requirement. For now, I will just address the question of whether the conditional probability account gives an adequate account of our ordinary concept of feasibility, but I will turn below to the thought that, although it does not do this, it provides an account of something that is a constraint on moral requirement (and to the thought that conditional probability premises, rather than feasibility premises, could complete the sorts of arguments for legitimacy that are the main subject of this thesis).

4.b)i) CPb rules out too much

Firstly, then, this specification of the account seems to concede *too much* to the agents' beliefs. Something may be feasible for an agent even though they do not believe there to be any effective means to it available. For instance, suppose that the only way for me to successfully run a 4-minute mile in the next month is for me to cut off my legs and replace them with enhanced 'super-legs' and that it would be very easy for me to do this. It may be that I do not believe this to be an effective means, perhaps just because it does not occur to me, or perhaps because I do not know that it is possible or that it would work. Nevertheless, there is at least some intuitive pull to say that it *is* feasible for me to run the 4-minute mile (at least in one sense) just because it *is* an option for me to cut off my legs and replace them with super-legs. There *also* seems to be a valid sense in which it is not feasible, but my account can allow for both of these senses (for one we hold fixed the agent's knowledge and for the other we do not), while the conditional probability account cannot. Similarly, I think intuitions are even clearer in a case where there are several apparent means available but only enough time to try one of them. Suppose I mistakenly believe that shouting at Geoff for a minute will make him angry and that this is the most effective means to make him angry; I also correctly believe that tickling Geoff for a minute will make him angry. Presumably the agent trying wholeheartedly will try the means she believes to be the *most* effective, so if I wholeheartedly tried to make Geoff angry in the next minute I would shout at him and I would fail. However, we intuitively want to say that it is feasible for me to make Geoff angry in a minute even though in the closest possible world in which I wholeheartedly tried I would fail. In this case, not only *is* there an effective means available, I am aware of it.

4.b)ii) CPb rules out too little

This version of the conditional probability account, then, rules out too much. On the other hand, it also seems to rule out too little: there are certain sorts of facts that at least sometimes seem to constrain feasibility claims that the conditional probability account cannot count as constraints. One straightforward case where this happens is the case of a sleeping or unconscious agent. Suppose I was fast asleep between 3 and 4 am. Was it feasible for me to call you at 3.30 am? There is a very natural sense in which the answer is no (it might be denied that *before* I went to sleep it was infeasible for me to call you at 3.30 am, but it is hard to deny that the sense is available in which at 3.15 am it is infeasible for me to call you at 3.30 am). Suppose, though, that I have all the correct beliefs about where the telephone is, how to use it to call you, and so on. *If* I had pursued the means I believed to be most effective (supposing, not implausibly, that beliefs can persist through unconsciousness), I would have had a high probability of success. The problem, though, is that, holding fixed my sleeping, it was not feasible for me to even try to pursue these means.[212]

4.b)iii) Motivational failure

The conditional probability account has further counterintuitive consequences in the case of motivational failure. There is an intuitive temptation to include at least some motivations as constraints on feasibility as well as some temptation not to always include all motivations as constraints.[213] In some cases we do not want the agent's motivations to count as a constraint on feasibility; that is, we want it to be feasible for A to bring about O despite the fact that A is not motivated to do so (such as in Brennan and Pettit's lazy parent case or Estlund's chicken-dancing case). In other cases, we want to allow an agent's motivations to count as a constraint on feasibility, such as in cases of pathological motivational failure. The conditional probability account cannot, I think, capture both of these intuitions where, again, my account can.

Account (CPb) can be read in two ways: 'pursue' in 'pursue whatever means to O are believed to be effective' can be read as a success verb or not. If read as a success verb, it means '*perform* (successfully) whatever actions are believed to be

---

[212] Lehrer (1976) 249 makes a similar point about conditional analyses of *'ability'*.
[213] Stemplowska (2016) argues that conditional accounts of feasibility fail adequately to deal with cases of motivational failure.

such that *if they were successfully performed* would be effective for O'.[214] If read in this way, the account seems to exclude all motivations, even extreme pathological ones. If I believe that walking across the plank positioned over the 500-metre chasm would be an effective means to cross it then there is an action that I believe to be effective that, if successfully performed, would have a high probability of resulting in my getting across the chasm. But if I suffer from a pathological fear of heights such that I could not bring myself to walk on the plank, I think we would be loath to say that it is feasible for me to cross the chasm. In any case, this is not a plausible reading of the condition, since it would make almost any *action* feasible for an agent: generally, one action I will believe to be effective for φ-ing is φ-ing, and if I successfully did that, then of course I would have a high probability of success in φ-ing.

If, on the other hand, we do not read 'pursuing' as a success term, we get something like '*setting out* to perform whatever actions are believed to be effective *if attempted*'. In this case, motivations seem to be ruled in as constraints more or less wholesale, since in the closest possible world in which I try wholeheartedly in *this* sense, it may be that I would not in fact succeed in performing the actions believed to be effective if attempted, just because I would not be motivated to do so: my attempt would be blocked by my motivations (or it could be that there are no actions believed to be effective if attempted because I know that I will not be motivated to carry through the attempt). This, too, seems implausible: we do not want to say that outcomes are infeasible for me whenever I lack the motivation to carry through actions/sequences of actions that would bring them about. A natural response is that 'pursuing' means neither 'performing' nor 'setting out to perform' but rather 'trying wholeheartedly to perform' but then the question is just postponed *ad infinitum*. On my account, these seemingly conflicting intuitions are accommodated since it simply says that on some sharpenings of 'feasibility' motivations are constraints, and on others they are not. If we hold fixed your motivations and there is no possibility of your intentionally performing an action to bring about your φ-ing given those motivations, then it will not be feasible for you to do so, but if we allow your motivations to vary and there is a possibility

---

[214] I simplify things here since, as can be seen, there are two possible readings of 'believed to be effective' and each of the two readings of 'pursue' could be matched with either of the two readings of 'believed to be effective'. The two possibilities not tried out here do not change matters.

(given the other constraints) of your φ-ing with different motivations then it will be feasible for you.[215]

Zofia Stemplowska defends a modified version of the conditional probability account that she thinks deals with the problem of motivational failure.[216] Her suggestion is that when there is a conceivable incentive 'that would bring the agent's motivational state in line with what is needed to perform the action in question', the action is feasible for the agent, whatever may be true about their *actual* motivations, but if there is no such conceivable incentive, then the action is infeasible for them. Her first proposal defines 'feasibility' thus:

> Action φ is (more) feasible if there is an incentive I such that, given I, X will try to do φ and, given I, is (more) likely to do φ.[217]

However, there is a problem with this, which is that there seems to be at least a sense in which it is *feasible* for me to, say, kill someone I love, even if there is *no* conceivable incentive that would induce me to do so. There may be few such actions, but we can certainly imagine there being some things that we are *so* motivationally committed to not doing that we never will (so long as our motivations remain constant). Nevertheless, it is natural to say that, in some cases at least, we are committed to not doing these actions *despite* their feasibility for us. Stemplowska notes this problem in the case of actions that we are committed to not performing for moral reasons: 'If [an agent's] failure to respond [to incentives] is solely due to her seeing action φ as (normatively) wrong, then we should not brand her as genuinely motivationally unable'. Thus, she revises the above definition to:

> Action φ is (more) feasible if there is an incentive I – or had the agent X not seen φ as wrong there would be I – such that, given I, X will try to φ and, given I, X is (more) likely to φ.[218]

---

[215] My account allows that it can be feasible for you to do something despite its not being possible for you to do it intentionally. For instance, supposing that holding fixed your motivations there is no possibility of you walking out onto a plank over a 500m chasm, it could still be feasible for you to walk out onto the plank if it is possible for you to press a button that activates a machine that takes control of your legs and makes you walk onto the plank. This may seem implausible, but I do not think that it is a serious problem. All that is needed is to distinguish between walking onto the plank, and the intentional action of walking onto the plank. In the above case, the outcome consisting in your walking onto the plank *will* be feasible for you, but the outcome consisting in your performing the intentional action of walking onto the plank *will not*. (If we understand the verb 'walking' as necessarily requiring intentionality, then there is no problem in the first place.) There is also, of course, the possibility of you taking a pill to change your motivations and thereby making the *intentional action* possible, but this, of course, is not consistent with your motivations being held fixed.
[216] Stemplowska (2016)
[217] Ibid. 280

However, this only resolves part of the problem. An agent's robust motivational commitment to not doing an action need not be for moral reasons. It does not seem inconceivable that you might be perfectly committed to not doing φ for non-moral reasons, but yet there be no (or little) motivational difficulty in doing φ *if you wanted* (and thus presumably no motivational inability). I might resolve to pursue some project (with no particular *moral* value) come what may and be so stubborn or determined that *no* incentive will induce me to do otherwise even though I *can* easily motivate myself to do otherwise *if* I want to. For this reason, I think my account does better than Stemplowska's: something we will never be motivated to bring about can be feasible when we do not hold motivations fixed.[219]

4.c) Mixing possibility and probability

Gilabert and Lawford-Smith present versions of the conditional-probability account that mix elements of possibility-based accounts with probability-based accounts.[220] They argue that there is both a binary *and* a scalar notion of feasibility. The binary sense is a matter of possibility (consistent with certain expansive 'hard constraints', such as logical, nomological and biological constraints), while the scalar sense is a matter of probability conditional on trying. The degree of feasibility in the scalar sense is supposed to be determined by 'soft constraints', such as economic, institutional and cultural constraints. In other words, the former 'hard constraints' are taken to make outcomes impossible (and thus rule them out as infeasible), while the 'soft constraints' merely make them less probable (and thereby reduce their degree of feasibility).[221]

I think they are right to note that there is a binary sense of 'feasibility', and it will be noticed that my account of feasibility given a choice of FC is similar to their account of binary feasibility, only replacing compatibility with hard constraints with compatibility with the chosen set of facts. However, I think it is wrong to think of constraints on feasibility as working in these two discrete ways. Social, cultural, and economic constraints can render things impossible. For example, *so long as the laws are as they are*, it is simply impossible for a non-citizen to become

---

[218] Ibid. 281

[219] I think she is right, though, that *when we hold motivations fixed*, actions that we may seem motivationally unable to do *are* feasible if there is some incentive that will induce us to do them.

[220] Gilabert (2009) and (2017); Lawford-Smith (2013); Gilabert and Lawford-Smith (2012). There are small differences between the accounts presented in these four papers, but the elements I discuss seem to be constant throughout.

[221] Tyler Cowen (2007) and Geoffrey Brennan (2013) have also argued that feasibility is (or at least can be) a matter of degree.

president of the USA. If the laws are fixed, there is no possible world in which a non-citizen becomes president of the USA; if a non-citizen claims to be president of the USA they will simply be wrong. This is of course a different kind of impossibility from that involved in contravening the laws of physics. The positive laws of the USA can be changed; the laws of physics cannot. (Even this distinction, though, does not obviously correspond to Gilabert and Lawford-Smith's distinction between hard and soft constraints: while laws of logic and physics can certainly not be changed, it may not be inconceivable for the laws of biology to be changed.) However, this difference is to do with whether the constraints themselves can be removed (how 'fixed', or hard to change, they are), not the manner in which the different kinds of constraints affect feasibility given that they are in place. The 'soft constraints' can affect the probability of outcomes, but they can also *rule out* outcomes. (They cannot, of course, rule out outcomes as, for example, physically impossible, but they can rule out outcomes as impossible so long as the soft constraint in question is constant.)

On the other hand, the laws of physics may render things improbable but not impossible. This point does not depend on the laws of physics themselves being probabilistic: a non-probabilistic law could contribute to some outcome's having a low probability. For instance, if we hold fixed the laws of physics as well as my knowledge, skills, strength and so on, I will have a low probability of surviving if I jump off a cliff into the sea. It is not *impossible* that I will survive, but if the laws of physics were different I could have a much higher chance of survival (for example, if the laws of physics were such that when an object falls towards earth it is repelled back towards where it started).[222]

Something like this problem motivated David Wiens to look for a binary sense of feasibility that is more restrictive than that suggested by Gilabert and Lawford-Smith.[223] He, like me, rejects the conditional probability account and offers a possibility-based account in its place, arguing that feasibility should be understood as possibility consistent with a 'resource stock': 'realising a target state of affairs is *feasible* only if there is an attainable resource stock that enables us to realise it'.[224] The resource stock defines an accessibility relation on the set of possible worlds and feasibility is a matter of possibility within this accessibility relation: in other words, there being a possible world consistent with the resource stock in which the outcome comes about. The accessibility relation is defined thus:

[222] Wiens (2015a) makes a similar objection to Lawford-Smith (450).
[223] Wiens (2015a)
[224] Ibid. 455

136

'a world is a member of the feasible set only if it is compatible with the facts pertaining to the composition of the total resource stock, the conversion processes, the causal processes and the state(s) of affairs that obtain at the actual world at the time of evaluation'.[225] Now, as I think Wiens is aware, it is not plausible to define feasibility simply in terms of possibility, even restricted possibility (he seems only to offer it as a *necessary* condition for feasibility). Such an account, if taken as definitional of 'feasibility', would, just like the simple possibility account, have implausible consequences of the sort raised by Brennan and Southwood. There is almost certainly a possible world consistent with the 'resource stock' in which our medical ignoramus successfully performs a neurological operation. In general, it would allow too many outcomes to count as feasible that one might succeed in bringing about only by luck. I think Wiens is right to reject the conditional probability view and to bring possibility back in, but I think it is also necessary to add an additional element to the simple possibility account, as my account does, to rule out outcomes that an agent could possibly bring about by luck.

This is not, though, something that Wiens need disagree with. Some additional elements could be added to Wiens's necessary condition to give a full account. His aim appears to be primarily to argue that there are binary constraints on feasibility (capable of ruling out proposals as infeasible) that are more restrictive than Gilabert and Lawford-Smith's hard constraints. However, even as only a necessary condition, Wiens's account has to rule out outcomes that could in some contexts reasonably be said to be feasible and it also has to count as feasible outcomes that could reasonably be said to be infeasible. Wiens defines his accessibility relation (the set of worlds that constitutes the feasible set) in two steps: it includes not only worlds that can be realised given the *actual* resource stock, but also worlds that are realisable given resource stocks that are *attainable* by transformation of the actual resource stock. *Either* we only allow *one* transformation of the resource stock (supposing there is some way of delimiting what counts as a *single* transformation) or we allow multiple iterations.

If we only allow worlds realisable after *one* iteration of transformation of the actual resource stock (whatever that means), then we seem to arbitrarily restrict the feasible set, and we rule out many things that we will intuitively want to count as feasible. For instance, suppose that to institute some policy we will need increased economic resources and in order to get these we will need to change

---

[225] Ibid. 457-8

public opinion about government spending (but this is quite easily done). This sounds like it will involve multiple successive transformations of the resource stock but it does seem that there is at least a sense in which it is feasible.

On the other hand, we could allow multiple iterations of transformation of the resource stock. This, though, seems to make Wiens's account of feasibility *very* permissive. There are *very many* (quite unrealistic) outcomes that come about in *some* possible world that is accessible by *some* possible series of transformations of our current resource stock. There is almost certainly, for instance, a *possible* series of transformations of the current resource stock that is consistent with the realisation of a proportional electoral system in the UK. But it does not seem obviously false to assert that the realisation of such a system is *not* feasible, in at least some valid sense. Wiens does note that his account has the consequence that it is very difficult to *know* whether distant outcomes *are* feasible.[226] This, though, is not the problem here. It is not natural to think that someone asserting the infeasibility of a proportional electoral system should rather be asserting simply that we do not *know* whether such a system is feasible. It in fact seems that we can have fairly good grounds for believing that such a system *will* come out as feasible on Wiens's account, but yet there seems to be a valid sense in which it is *not*. Any attempt to give a single set of necessary and sufficient conditions for feasibility will, I think, miss the fact that for many outcomes there is both a sense in which they are feasible and a sense in which they are not. A proportional electoral system in the UK is feasible in one sense in that it is perfectly consistent with the deeper facts of human society and so on, but plausibly is not feasible in another sense, which holds fixed more of the actual facts.

4.d) The constraint on moral requirement

I have argued in this section that available alternatives to my multivocal account fail to capture the concept of feasibility. Proponents of these alternatives, however, might grant that their accounts do not provide an analysis of the ordinary concept of feasibility, but claim that they *do* provide an account of some other concept that is a constraint on moral requirement. Thus, for instance, it might be claimed that, even if feasibility is not conditional probability, conditional probability is a constraint on moral requirement: something like an "'ought" implies "conditionally probable"' principle holds. Even if I have found the best analysis of the ordinary concept of feasibility, it might be said, this just shows that

---

[226] Ibid. 467

feasibility is not the constraint on morality that it was thought to be; the constraint is rather conditional probability (or something else).

The reason that we are interested in feasibility here is that plausible moral arguments for state legitimacy appear to depend on premises about feasibility. It is standardly thought that feasibility has some relevance to morality (it is usually thought to be a simple *constraint*). If it is morally relevant, then it makes sense to expect to be able to make arguments about moral permissibility using premises about feasibility. However, if you think that it is instead something else (like conditional probability) that is morally relevant, then you might think that the arguments for state legitimacy can be made instead using premises about conditional probability (or whatever else is the constraint on moral requirement). In that case, my account of feasibility will have no particular bearing on state legitimacy.

The plausible arguments for state legitimacy that I discussed in the first part of the thesis required premises like 'This state's being an exclusive enforcer is the only feasible way of doing something morally required (e.g. achieving full status freedom) or morally good' or 'The best feasible way things might be involves this state's being an exclusive enforcer'. If feasibility is a constraint on moral requirement (that is, an action can only be morally required if it is feasible), then it seems plausible that, for instance, if an action is the only feasible way of doing something morally required, then it too is morally required. If a) you are morally required to φ, b) only feasible actions can be required and c) the only feasible way to φ is to ψ, then plausibly you are required to ψ. (It might also be thought to follow that an action which makes things the best they might feasibly be is permissible, though that is more controversial.) However, if it is not feasibility that constrains morality in this way, but, say, conditional probability (that is, an action can only be morally required if it is probable conditional on trying), then a similar principle for conditional probability might plausibly follow. In that case, it will be possible to make versions of the arguments for state legitimacy discussed above based on conditional probability premises (for instance). Instead of a feasibility premise, we might have something like the following: 'this state's being an exclusive enforcer is the only way of doing something morally required that has a sufficient probability of success conditional on trying'. If conditional probability is a constraint on moral requirement, then this might be sufficient to establish permissibility.

However, I think that not only do the alternative accounts discussed in this section fail to give a plausible account of feasibility, they also do not identify a straightforward constraint on moral requirement. In fact, I think, the arguments I gave above against these alternative accounts as accounts of feasibility also show that they fail as accounts of a constraint on moral requirement. I will briefly run through them in turn to indicate how they transfer.

Firstly, it seems plausible that you can be morally required to φ even though in the closest possible world in which you pursue the means you believe to be effective for φ-ing, you would be unlikely to succeed. In my example above, I mistakenly believe that shouting at Geoff for a minute will make him angry and that this is the best way to do so, but I correctly believe that tickling him for a minute will make him angry. In the closest possible world in which I pursue the means I believe to be effective for making Geoff angry in the next minute, I presumably shout at him for a minute and fail. It seems plausible, though, that this is compatible with my being morally required to make Geoff angry in the next minute in this case (leaving aside *why* this might ever be morally required of me). To take another case, it seems plausible that Daphne fails to do something she was morally required to do if she votes for a racist party A, believing it to be an anti-racist party. Whether she is *culpable* for this moral failing will presumably depend on whether she was culpably ignorant. But suppose she was culpably ignorant. If (CPb) is a constraint on moral requirement, then we cannot say that Daphne failed to do something she was morally required to do (vote for a non-racist party) because she was culpably ignorant. In the closest possible world in which she pursues the means she believes to be effective for voting for an anti-racist party, she fails. Nevertheless, it is natural to think that she *did* fail to do something she was required to do. If we take (CPb) as a constraint on moral requirement, we can still say that she was required to inform herself about the parties, and that she failed to do *this* morally required thing, but we cannot say that this led her to do another wrong thing (vote for the racist party).

Secondly, as we saw above, if we read 'pursue' in (CPb) as a success verb, almost any action turns out to be conditionally probable, rendering the proposed constraint on moral requirement bloodless. On the other hand, if we do not read it as a success verb, we get something like the claim that required actions must be likely to succeed conditional on '*setting out* to perform whatever actions are believed to be effective *if attempted*'. This makes motivations a constraint on moral requirement pretty much wholesale. If you lack the motivation to carry through

an attempt to φ, then you will be unlikely to succeed if you set out to perform the actions you believe to be effective for φ-ing if attempted. It must be possible, though, to be morally required to do things that you lack the motivation to do.[227] If it is possible to be morally required to do something even though you would not be likely to succeed given wholehearted trying, then something's (ψ) being the only way of doing something morally required (φ) that would be likely to succeed conditional on wholehearted trying does not show that it too is morally required (since you could be required to φ in some other way).

It might be thought that (CPa) (probability conditional on performance of the *objectively best* available action), which we rejected as an account of feasibility, provides a constraint on moral requirement. It obviously cannot be *the only* constraint on moral requirement, since the medical ignoramus who performs the objectively best sequence of actions will be likely to successfully perform a neurological operation, yet it is not plausible that they could be morally required to do so. But it might be thought that (CPa) is nevertheless *a* constraint on moral requirement: it is *necessary* that something be probable conditional on the objectively best action for it to be morally required. This claim, though, adds little to the thought that you cannot be required to do what is not feasible for you. The notion of *availability* it makes use of is just as much in need of elucidation as that of feasibility. We could understand it in terms of conditional probability, but then all of the above problems resurface. If instead we understand it in terms of my account of feasibility, or in terms of possibility, we end up with multiple possible ways in which the notion could be made precise, and so it is unclear, in just the same way as for feasibility, exactly what constraints this does put on morality.

Stemplowska's definition of feasibility also cannot be a constraint on moral requirement. She proposed the following:

> Action φ is (more) feasible if there is an incentive I – or had the agent X not seen φ as wrong there would be I – such that, given I, X will try to φ and, given I, X is (more) likely to φ.[228]

As I argued above, it seems perfectly possible that one could be so committed to not doing something for non-moral reasons that there is no incentive that will in fact lead you to try to do it. And it also seems perfectly possible that one could be morally required to do that thing despite this extreme level of commitment. If somebody has such an extreme commitment, for example, to not touching people

---

[227] See Estlund (2011).
[228] Stemplowska (2016) 281

of some gender that no incentive could induce them to do so, this does not rule out their being morally required to save someone of that gender from drowning (even when this requires physical contact).

Gilabert and Lawford-Smith's 'hard constraints' could act as a constraint on moral permissibility (and they are designed to do so). However, as we saw, they are not the only constraints that, insofar as they are held fixed, can rule out proposals. It seems somewhat arbitrary to draw the line exactly where they have done. Perhaps it is true that the constraints they point to (logical, nomological and biological constraints) are always, or almost always, constraints on moral requirement. But it does seem plausible that these are not in all contexts the *only* feasibility constraints. Their 'soft constraints' on the other hand cannot play a constraining role, and they are not designed to, since they are scalar. There could be some threshold of scalar feasibility below which actions cannot be morally required, but if we say this we just end up with the standard conditional probability account again, since that account just restricts moral requirement to those actions that are *sufficiently* probable conditional on trying (i.e. those above some threshold of conditional probability). We saw above that this does not identify a general constraint on moral requirement.

Wiens's account of feasibility is an attempt to argue that there is a binary feasibility constraint on moral requirement that is stronger than Gilabert and Lawford-Smith's hard constraints. But, as I argued above, it seems either to arbitrarily restrict the feasible set in a conservative way to what can be achieved through one iteration of transformation of the resource stock, or else to be excessively permissive. On the one hand, it does not seem plausible that we can never be morally required to pursue outcomes that require reasonably long-term planning, and multiple iterations of transformation of the existing resource stock. On the other, there are very many unrealistic outcomes that come about as the result of *some* series of possible transformations of the current resource stock. It seems quite plausible that there are constraints on moral requirement more demanding than this.

Thus, I do not think that these alternative accounts provide a straightforward constraint on moral requirement. If something is feasible for you (on some relevant sharpening) but not conditionally probable (that is, would not be likely to come about if you pursued the means you believe to be effective), it may still be morally required. If it is *possible* for you to do it or bring it about reliably given certain aspects of the way the world is (i.e., holding them fixed), you may be

morally required to do it even if you would not be likely to succeed in the closest possible world in which you pursue the means you believe to be effective. It does not matter what happens in the closest possible world in which you try wholeheartedly; what matters is what is *possible* for you (but the question remains, possible given *what?*). It seems natural to think that the constraints on what you can be morally required to do are not given by what would happen in some particular counterfactual possible world, but by what *is* possible for you given certain facts of the actual world. Consequently, there is no reason to suppose that the arguments for state legitimacy I discussed in the first part of the thesis go through if made with, for example, conditional probability premises instead of feasibility premises. If, for instance, it is the case that being an exclusive enforcer is the only way for state S to ensure status freedom (which, let us suppose, is morally required) that is likely to succeed conditional on wholehearted trying, it does not follow that it is permissible, since there could be other ways of doing it that are morally required despite not being likely to succeed conditional on wholehearted trying.

My claim is not that feasibility *is* a straightforward constraint on moral requirement, and these other concepts are not. As I will argue in the next chapter, my account of feasibility has the consequence, on the contrary, that feasibility *as such* cannot be. There is no single sharpening of 'feasibility' that is obviously privileged over others, and most outcomes are feasible on some sharpenings and not on others. It is not obvious *which* facts an action must be compatible with in order to be morally required. Feasibility tout court, then, is not apt to be a straightforward constraint on morality: at best *some* sharpening(s) of feasibility can be (and most likely, different sharpenings are in different contexts). But I think that whatever constraints there are on moral requirement (though they may not be as simple as they are often thought to be) are most plausibly to do with feasibility (or possibility), not conditional probability or other alternatives. Since none of these, though, provide a straightforward constraint on moral requirement, arguments for state legitimacy that rely on feasibility premises (or conditional probability, or similar) cannot be taken for granted. Whether they are read as being about feasibility as defined by my account, or about conditional probability, or any other alternative discussed above, they cannot be assumed to make the arguments for legitimacy go through.

5.   A multivocal conditional probability account?

The above arguments suggest that my multivocal account does better at capturing the concept of feasibility than alternative available univocal accounts, and also that these latter accounts do not identify a constraint on moral requirement (and so do not offer a straightforward alternative way of arguing for state legitimacy). However, one might accept these arguments and this conclusion, and accept the need for a multivocal account of feasibility, but think that the correct multivocal account should be cashed out in terms of *conditional probability*, rather than in terms of possibility, as in my account. I will call such an account, snappily, the *multivocal conditional probability* account (MCP). Such an account would accept my primary contention that the best account of feasibility we can give is a multivocal one, and so would have similar consequences for arguments for state legitimacy to those that follow from my account (which I will discuss in the next chapter). Still, I will briefly argue here that there is no obvious way to make an MCP account work, and so we should take my possibility-based account to be the best going attempt at a multivocal account of feasibility.

We could give such an account as follows. When we ask about the feasibility of φ-ing for me in the actual world, we say (just as I do) that there are multiple possible ways of making this precise, and we will only get a determinate answer once we specify a sharpening. On this account, though, a sharpening is just a specification of the facts of the world; call it a *starting point* (SP). Unlike for my account, we do not ask which of the facts that actually hold we will hold fixed, rather we just choose a set of facts. These could be any facts: while normally we will be interested in starting points that have some resemblance to the actual world, we simply choose the facts however we like. There will then need to be some account of how the facts chosen on some SP must relate to the facts of the actual world in order for it to count as a possible sharpening of feasibility for *me* (or any actual agent). An MCP account *must* do something like this rather than choosing a set of the *actual* facts to hold fixed (like my account) as I will explain below.

We can then define (agent-relative) feasibility on an SP thus:

(MCP) It is feasible for agent X to bring about O on SP *s* iff, with the facts being as specified by *s*, if X tries wholeheartedly to bring about O, X has a reasonable probability of success.[229]

I do not think, though, that the account is viable, and for this reason think that we should consider the possibility-based account to be the best available multivocal account.

There is a basic problem with the MCP account, which is that there is no obvious way to determine which sets of facts relate to the actual world in the right way for an SP to count as a sharpening of 'feasibility' for *me* or *us*. My account *starts* with the facts of the actual world, and then defines sharpenings as choices of which of *those* facts to hold fixed. Admittedly, it allows for outlandish limiting cases, where barely any of the facts of the world are held fixed, but these are still limiting cases of feasibility for the actual world; an FC could not hold fixed any facts that *do not* hold in the actual world. The MCP cannot do things this way. If we choose an arbitrary set of facts of the actual world to hold fixed and then ask what would be likely to happen if we tried to φ, the choice of facts to hold fixed makes little difference to the answer. For instance, even if we allow, say, all of the actual facts to vary except the laws of physics, the truth of a conditional probability statement will still depend only on nearby possible worlds in which the agent tries wholeheartedly. Even if we allow to vary all but the laws of physics, it will still be the case that if I tried wholeheartedly to run a mile in 5 minutes I would fail, because there is a closer possible world in which I have human legs and limited strength and so on than any in which I do not. Thus, for the multiple sharpenings to make a significant difference, they need to give us a world from which conditional probability is to be evaluated, not just a set of actual facts to hold fixed. For instance, if we assumed a starting point where the laws of physics were different, then it might be that if I tried wholeheartedly to run a mile in 10 seconds I would be likely to succeed. Or if we assumed a starting point where people were motivated differently, we might be likely to succeed in establishing some social structure that we would be unlikely to successfully establish in the actual world.

It is plausible that when we make a feasibility claim we tacitly assume a choice of which facts of the world to hold fixed and then mean something to do with what is possible given these facts. It is not plausible, though, that, as the MCP account

---

[229] We could modify the conditional probability part of this definition in the ways suggested above or in the ways suggested in Stemplowska (2016) to deal with the problems that motivated these modifications.

would have it, we are making a claim about what counterfactuals *would* hold in some *specific* non-actual possible world. On this view, feasibility claims that appear to be about the actual world lose their connection to it. There is no obvious principled way to determine which SPs (possible ways the world might be) can give sharpenings of claims about feasibility for *us* in the actual world. The MCP account also, relatedly, has the consequence that in order to make a determinate feasibility claim (with sharpening specified), I need to know *how* the world must be for the agent to bring about the outcome in question. On my account, when we make a feasibility claim, we just assert that there is *some* possible world out of a range in which the agent brings about the outcome, but on the MCP account, we assert that in a *specific* possible world (the closest world to the world specified by the SP in which the agent tries to bring about the outcome) the agent succeeds. This seems implausible. Thus, I think there is reason to prefer my account to the MCP account.

6. Conclusion

In this chapter, I have attempted to understand what we mean when we say that some outcome is or is not feasible. I have argued that no single binary definition is obviously privileged as the definition of feasibility *tout court*; rather, there is a whole range of possible binary 'sharpenings', corresponding to FCs, or selections of facts to hold fixed. Which of these is meant by some feasibility claim may be determined by the context or background assumptions. Alternatively, it may be indeterminate, requiring further specification for the claim to have determinate truth conditions. In principle, what I have said does not rule out the possibility of someone giving an argument that one of these sharpenings is in some way privileged over the others, but unless such an argument emerges, I think there is no reason to think that feasibility is a univocal concept. This, I believe, captures our ordinary concept of feasibility better than any of the rival accounts. There are many different ways in which we use 'feasibility', many different constraints that are tacitly assumed when we make feasibility claims. When we say, for example, that a proportional electoral system is not currently feasible in the UK we usually tacitly hold fixed a different range of facts to when we say that communism is made infeasible by human nature. I argued that a unified account of feasibility given a choice of FC can be given, illustrating how the different sharpenings are all sharpenings of a single concept. Feasibility, I claimed, is a matter of possibility consistent with the facts held fixed by the FC. This is the most plausible multivocal account of the concept available. An FC can thus be understood as functioning like an accessibility relation on a domain of possible worlds.

As I will argue in the next chapter, this account of feasibility has the consequence that we cannot take the success of arguments that depend on feasibility premises for granted until we have some idea of exactly which sharpenings of 'feasibility' these premises must be true on to make the arguments work. I also argued in this chapter that the alternative available accounts of feasibility, in addition to failing as accounts of feasibility, do not identify a straightforward constraint on moral requirement, and there is no reason to think that in considering the consequences of my account of *feasibility* for arguments for legitimacy, we are focusing on the wrong concept. If anything, I think, it is *feasibility* that will allow us to argue for the permissibility of some action, but there is no immediate or obvious answer to the question which sharpenings of feasibility are relevant as constraints for which moral claims.

Chapter 6:
Conclusions: unrealistic anarchism

Let me begin by taking stock of what we have done so far. The conclusion of the first part of my thesis was that the only arguments that could plausibly establish a general and exclusive permission to enforce for at least some existing states depend on feasibility premises. I argued that attempts to show that the legitimacy of certain kinds of states can be established *a priori* are unsuccessful, and the most plausible grounds we have for thinking that the legitimacy of at least some existing states is on safe ground depend on assumptions about feasibility. If there were a simple '"ought" implies "feasibility"' constraint on morality, and feasibility were a simple binary and univocal concept, then, as we saw, some feasibility premises like the following might well be sufficient to establish the legitimacy of some state S (that is, its having a general and exclusive permission to enforce):

> 'The best feasible worlds involve state S being an exclusive enforcer'.

> 'A condition of right can only feasibly be achieved or maintained with the existence of exclusive enforcer S.'

> 'Full status freedom can only feasibly be achieved or maintained with the existence of exclusive enforcer S.'

> Etc.

(Whether or not these premises *would* establish legitimacy I leave open, but it does not seem implausible that some such premise, if true, would suffice.)

In general, it is plausible that something like the following principle is true:

> (P1)   If x is the only way for agent A to do something A is morally required to do, then A will be morally required to do x, and so, a fortiori, permitted to do it.

In addition, though much more open to doubt, it could also be thought that another similar principle holds:

> (P1')   If x is the only way of bringing about the best possible state of affairs, then it is at least permissible to do it.

Certainly, on a consequentialist view, P1' will hold. If there is a deontological complaint against x, then P1' need not hold, but it might be thought that it does generally hold where the complaint against x is only pro tanto, i.e., is defeasible.

If, then, it is morally required, for instance, to create and enter, or maintain, a condition of right, then if state S's being an exclusive and general enforcer is the *only* way this could be done, it follows from P1 that it has an exclusive and general permission to enforce. The claim, though, of the first part of the thesis is that there are no plausible strict necessity claims of this sort capable of making arguments for state legitimacy go through. The most plausible way to make these arguments work involves *feasibility* premises instead of strict possibility premises. It might seem plausible, though, that, as mentioned above, as well as P1 (and perhaps P1'), a similar principle involving feasibility is true:

> (P2)   If x is the only *feasible* way for agent A to do something A is morally required to do, then A will be morally required to do x, and so, a fortiori, permitted to do it.

And again, more controversially, perhaps also:

> (P2')   If x is the only *feasible* way of bringing about the best possible state of affairs, then it is at least permissible to do it.

(We saw in the last chapter that there are certain alternative candidates to feasibility premises to complete arguments for state legitimacy, such as conditional probability premises, but we saw also that it is not plausible that an equivalent of P2 (or of P2') holds for these.)

Principle P2 (and perhaps P2') seems plausibly to be the corollary of an '"ought" implies "feasibility"' principle. If you can only be morally required to do the feasible, then your being morally required to do something must be your being morally required to do it in a feasible way, and if there is only one feasible way to do something morally required, then it must be morally required to do it in *that* way. If P2 holds, then, and if feasibility is a simple binary and univocal concept, we can see how premises such as those mentioned above will suffice to establish the legitimacy of a given state. If it is morally required to create and enter a condition of right, for example, and the only feasible way of doing so involves state S's being an exclusive and general enforcer, then, by P2, it must be permissible for A to be an exclusive and general enforcer.

However, the last chapter has the consequence that matters are not quite so simple. I argued that the best account of the concept of feasibility is a multivocal

one; there are a good many different possible binary sharpenings of the concept, no single one of which is obviously privileged in general or for the purposes of moral theory. Thus, if there is an '"ought" implies "feasibility"' principle with which moral and political theory must deal, it is not a straightforward matter what it is. Perhaps there is some single sharpening of 'feasibility' of the many my account alludes to which is the constraint on moral 'oughts', but it is not obvious which one it is, and nor is there any immediate reason to suppose that there is a single one. It could well be that different feasibility constraints are relevant to different sorts of moral claim.

Thus, whether or not a principle like P2 is true is not a straightforward matter. It is certainly not true for all sharpenings of 'feasible' (all FCs). For instance, if we filled in P2 with a sharpening of 'feasible' that holds fixed *all* of the facts of the world (the most restrictive FC), then it would have us being permitted to do x just because it is the method that we are currently using to do something morally required. The only feasible way of doing something morally required on such a restrictive understanding of feasibility is the way that it is currently being done. It certainly does not follow, though, from the fact that x is the method we are currently using to do something morally required that x is permissible. Unless we are consequentialists, we should be open to the possibility that there are *impermissible* ways of doing morally required things. There are thus at least some sharpenings of 'feasible' for which a P2-like principle does not hold. Plausibly it is not just the *most* restrictive FC that is too restrictive for such a principle: for at least a number of different sets of the actual facts, something's being the only way of doing something morally required that is consistent with holding fixed that set of facts does not entail its permissibility. Thus, it seems that principle P2 will be true given some sharpenings of 'feasible' (some FCs), but not others (i.e., it may be true if you replace 'feasible' with 'feasible-on-FC-f' for some FCs, but not for others).

In addition, it may well be that no version of P2 holds uniformly for all possible actions, but that different versions of P2 (involving different sharpenings of 'feasible') are true for different possible actions. (That is, it may well be that the sharpening on which it is true that x's being the only feasible way to do something morally required is sufficient for x to be morally permissible is different from the sharpening on which it is true that y's being the only feasible way to do something morally required is sufficient for y to be morally permissible.)

Thus, without an account of which sharpenings (which FCs) make P2 (or some similar principle) true (which will presumably be closely related to an account of which FCs make an '"ought" implies "feasibility"' principle true), we cannot assume that the feasibility-dependent arguments for states' legitimacy go through. It will need to be the case that the relevant feasibility premises are true on the same sharpening(s) that make P2 true (or make it true for the relevant subdomain of morality). If the sharpening(s) of feasibility that makes P2 true is *not* one on which, for example, state S's being an exclusive enforcer is the only feasible way of achieving a condition of right, then these arguments will not be valid. Thus, it cannot be assumed that the legitimacy of some existing states is on safe ground unless there is some good reason to think that the sharpening of 'feasible' on which P2 is true (for exclusive enforcement) is one on which the only feasible way of achieving a condition of right (or whatever else) does involve these states having exclusive enforcement power.

One might have reason to think this if one thinks that the sharpening of feasibility that makes P2 true, and the sharpening that constrains moral requirement or 'ought' claims in general, is quite a restrictive one. We saw already that the *most* restrictive FC (which holds fixed *all* of the actual facts, and so makes nothing but the status quo feasible) does not plausibly make P2 true. While presumably nobody would think that morality is constrained by *such* a restrictive feasibility constraint, it could be thought that it is constrained by a *fairly* restrictive FC, which holds fixed *many* or *most* of the actual facts. If φ's being morally required entails its being feasible on some restrictive FC *f*, then if φ *is* morally required and the only way of φ-ing that is feasible on *f* is ψ-ing, then ψ-ing must be morally required. If status freedom (to take one example) being achieved without the state's exclusive enforcement is a somewhat *distant* possibility, then morality's being constrained by a *restrictive* FC will have the consequence that the requirement to achieve status freedom necessitates the state's exclusive enforcement (as the only way of meeting the requirement that is feasible on the restrictive FC), and thus makes it permissible.

I will thus attempt to argue in this chapter that there is little reason to suppose that morality must generally be constrained by a restrictive sharpening of feasibility. I will first borrow from discussion by David Wiens (which in turn draws on work by David Estlund) to suggest that moral principles are relative to possibilities or feasibility facts. I will argue that my account of feasibility creates significant uncertainty about what the set of possibilities *is* that moral principles

for a given factual context are relative to. This in itself creates doubt about whether the legitimacy of some existing states is on safe ground: if the moral principles for the actual factual context are relative to a set of feasibility facts that do not make the necessary feasibility facts true, then we have no argument for the legitimacy of some existing states. Thus, an argument would be needed that this is not the case.

I then suggest that the possibilities relative to which it is *worthwhile* identifying principles need not be just those that give us the sets of principles that are in fact true in our factual context. I argue, further, that theory identifying relatively *unrealistic* principles (i.e., constrained by non-restrictive feasibility constraints, which hold fixed relatively few of the actual facts) can be worthwhile and useful for guiding action in the real world. My conclusion will thus have two parts. First, we cannot be confident that state legitimacy is on safe ground, since it is not obvious that the sharpening of feasibility that makes a principle like P2 true is one on which the feasibility premises necessary to establish state legitimacy are true. Second, even if these premises do come out as true on the relevant sharpening (and so some existing states *do* in fact have a general and exclusive permission to enforce), it is still worth knowing that these states' enforcement is *not* permissible relative to less restrictive feasibility constraints. I will conclude by exploring how this knowledge can be worthwhile and what implications it has for us in the actual world.

1. A 'realisticness' scale

Feasibility constraints (FCs), as I have defined them, are sets of facts. An FC is the set of facts that is held fixed and relative to which feasibility is assessed by the corresponding sharpening of 'feasible'. We can make sense of the idea of a scale along which these FCs are ranked in terms of 'realisticness'. The most 'realistic' or 'restrictive' FC, the lowest on the scale, is the one mentioned above that holds fixed *all* of the facts of the given factual context. At the other extreme, the most *unrealistic* FC, the highest on the scale, holds fixed *no* facts (except perhaps the laws of logic). Between these two extremes, we progress up the scale as we allow to vary more and more facts, and facts that are more and more 'unchangeable'. On the lowest, most restrictive FC, the only thing that will come out as feasible is the status quo. On the highest, least restrictive FC everything will come out as feasible (except perhaps the logically impossible). In general, as we progress up the scale, more and more outcomes will come out as feasible.

It is a complex and difficult matter what exactly determines these rankings and we do not need a full answer. What will determine the position of an FC in the ranking is a measure of how *plausible* is a possible world in which all of the facts that the FC allows to vary (all of those that it does not hold fixed) do not hold. What fixes the plausibility of possible worlds is not a question I shall attempt to answer, but it will suffice to notice that we do often have a very rough intuitive ability to judge this. While we are certainly unlikely to be able to make a full ranking with any confidence (and it may not be *possible* to make a full ranking), there are a number of pairwise comparisons about which we will have no hesitation. Thus, although without a measure of plausibility of worlds we cannot suppose that the ranking of FCs on the realisticness scale will be *complete*, we can safely assume that we will be able to achieve a rough ranking of at least a number of FCs. It will be clear for a number of FCs that they are either *quite realistic* or *quite unrealistic*. Given such a scale, then, we can understand the idea that might support an assumption that the legitimacy of some states is on safe ground as the idea that morality is generally constrained by a *quite realistic* FC. It may well be the case that the only way of meeting a moral requirement like the requirement to create and enter a condition of right that is feasible on a *realistic* or restrictive FC involves the state having general and exclusive enforcement power. Thus, if P2 or a similar principle is true for such a realistic understanding of 'feasible', it may be possible to make a successful argument for the legitimacy of such a state. My aim now, then, will be to argue that it is not *obvious* that morality is constrained by a sufficiently realistic FC and that moral theory constrained only by quite *unrealistic* FCs can be worthwhile.

2. *Possibility-relative principles*

David Wiens has recently defended a model for understanding normative principles that begins from what he calls the 'Uncontroversial Thesis'.[230] He divides moral principles into 'evaluative principles' (which 'serve to comparatively assess and rank options according to some set of normative criteria') and 'directive principles' (which 'perform a deontic function' by marking 'the lines between obligatory, permissible, and impermissible options').[231] The 'Uncontroversial Thesis' states that '*a set of directive principles is justified relative to a particular set of salient possibilities*', and, more precisely, 'which options – actions, institutional schemes, states of affairs – an agent is required, permitted, or prohibited to realise depends on the set of options that is open to that agent in

---

[230] Wiens (2017)
[231] Ibid. 152-3

153

some (for now, unspecified) sense'.[232] The idea is that what an agent can be *required* to do is not insensitive to what it is *possible* for the agent to do. The thought is similar to the common thought that 'ought' implies 'can'.

Wiens presents a number of pieces of evidence to support the claim that this thesis holds (and holds uncontroversially). These include the fact that much of the debate concerning the use of idealisations in normative theory depends on the idea that assumptions about the set of feasible options matter for the content of the normative theory, as well as the fact that the standard semantics for deontic modals fits with the thesis. One important thing (for our purposes) he points to is Jackson and Pargetter's much-discussed case of Professor Procrastinate.[233] Professor Procrastinate is given a review assignment, but knows that he is disposed to procrastinate, and will not in fact complete the review on time if he accepts it. The standard judgment is that he ought to [accept *and* complete] the review, but given that he will not complete it, he ought not, all-things-considered, to accept the review. Relative to a set of possible worlds that includes one at which he completes the review, he ought to accept it. But relative to a set of worlds at which he will fail to complete the review, he ought not to accept it. Thus, it seems, as Wiens concludes, that 'the directive principle to which Procrastinate is subject is sensitive to the set of possibilities one deems salient'.[234]

The 'concessive' directive principle (to use Estlund's term), i.e., the principle requiring Procrastinate to refuse the review given that he will not complete, disappears when the salient set of possibilities includes worlds at which he does complete the review, and thus at which the 'non-concessive' principle is satisfied.[235] The non-concessive principle, though, does not disappear when the salient set of possibilities only includes worlds at which he does not complete the review, i.e., worlds at which the non-concessive principle is not met. Even if he will not complete, and so ought not to accept, it remains the case that he *ought* to [accept *and* complete].

Estlund argues that this asymmetry shows that the non-concessive principles have a certain kind of *primacy* over the concessive ones.[236] In order to accommodate this appearance, Wiens suggests what he calls a 'nesting model' of directive principles. According to this model, one directive principle appears to

---

[232] Ibid. 154
[233] Jackson and Pargetter (1986)
[234] Wiens (2017) 154
[235] Estlund (2017a) 38
[236] Ibid. 39-40

have primacy over another because the second is specified relative to a set of possibilities that is a subset of the set of possibilities relative to which the former is specified. If a principle A is relative to a set of possibilities that *includes* all the possibilities that another principle B is relative to, then A will appear to have some kind of *primacy* over B. Wiens thus explains the Professor Procrastinate example as follows:

> Let *C* be the set of whatever facts about Procrastinate determine that he will not complete this particular review assignment … and let *D* be the set of background facts that are salient to the case less those in *C* … Now let *V* be the set of worlds that are consistent with the facts in the union of *C* and *D* and let *W* be the set of worlds that are consistent with the facts in *D*. Notice that *V* is a proper subset of *W* … Now we restate the standard pair of judgments: Procrastinate ought to refuse the assignment given *V* because he refuses at the highest-ranked worlds at which he fails to complete; he ought to accept (and complete) the assignment given *W* because he accepts and completes at the highest-ranked worlds in *W*. Put this way, we can see how natural it is to think that the former directive "evaporates", giving the latter directive some sort of primacy. The (concessive) directive to refuse is specified relative to *V* and, thus, only obtains at the worlds in *V*, whereas the (nonconcessive) directive to accept and complete is specified relative to *W* and, thus, obtains at the worlds in *W, including the worlds in V.*[237]

This explanation of the case draws on Wiens's optimisation model of normative theory, according to which identifying directive principles is a matter of 'optimisation' of rankings of possible worlds within a feasibility constraint. Wiens denies that his optimisation model relies on a consequentialist picture of morality, but it is not quite clear to me that we really can think of deontological principles as establishing a ranking of possible worlds (rather than just a *partition* of possible worlds) without losing what makes them distinctively *deontological*. I think, though, that the nesting model described above need not depend on the optimisation picture. We need not say that a directive principle's holding relative to a set of possible worlds is a matter of its being realised at the *highest ranked* worlds in that set. Rather we can just say that a *necessary* condition on a directive principle's holding relative to a set of possible worlds is its being realised at some member of that set. This is just to state a version of the "'ought' implies 'can'" principle. Whatever the salient set of possibilities is, a directive principle must not require something that is impossible given the salient set.

Further, it seems that a directive principle's holding relative to a set of possibilities requires not only that there be some world in the set at which what it requires comes about, but that what it requires be *feasible* given the set of possible worlds in question (or given an FC). On my account of feasibility, an outcome's

---

[237] Wiens (2017) 162

being feasible for an agent given an FC requires not only that it *come about* at some world in the set fixed by the FC, but also that there is a world in the set at which the agent performs an intentional action that brings about the outcome safely and competently. It seems irrelevant to a principle's holding relative to a set of possibilities that what it requires comes about at some world in that set if it only comes about by luck or not as a result of intentional action. Thus, I think the necessary condition for a principle to hold relative to a set of possibilities is that it be *feasible* given the set of possibilities. Procrastinate ought to accept and complete the assignment given W because there is some member of W at which he does so as the result of intentional action, safely and competently; it is feasible for him to do so. However, there is no member of V at which he does both, so given V he ought to refuse, because if he will not complete, he ought not to accept. Still, because all of the worlds that are members of V are also members of W, it remains the case in any world at which he will not complete that he ought to [accept and complete].

I will return below to the idea of primacy and hierarchies of principles, but for now let me make use of the idea that directive principles are relative to sets of possible worlds, and the idea that it is a necessary condition on a principle's holding relative to a set of possible worlds that it be feasible given that set. That directive principles are relative to sets of possibilities means that the directive principles that hold vary across different factual contexts, since presumably what is *possible*, or what is *feasible* (i.e., what the salient set of possible worlds is) varies with factual context. Wiens describes his 'Uncontroversial Thesis' as the thesis that what an agent is required or permitted to do depends on the options that are open to that agent. Any plausible way of cashing out what the set of options open to an agent is will presumably need to be able to deliver the result that these options vary with factual context.

There is, though, a question as to *what* the salient set of possibilities is for any *given* factual context. Wiens takes the set of options to be given by a set of *feasibility constraints*.[238] The account I have given of the concept of feasibility, though, has the consequence that for any given factual context there is no single determinate set of options that are the feasible options (for a given agent). There is no single determinate set of possible worlds that are relevant to an outcome's feasibility. For different FCs, different sharpenings of 'feasible', different principles will be ruled out. Thus, for a given factual context, there are many

---

[238] Ibid. 157

different sets of feasible options and so, many different possible sets of principles whose realisation is feasible. The idea that the realisation of a directive principle must be feasible does not fix a single salient set of possibilities for a given factual context. It might be that there *is* some particular sharpening of 'feasibility' that is the salient one for directive principles, but given my account of feasibility, there is no particular reason to suppose that there is. An argument would be needed to show that there is some such salient sharpening. Wiens himself argues for a particular binary feasibility constraint on political theory, but as argued above, it seems either to be too restrictive or too permissive (depending on how it is read). It is not plausible to think that the constraint he proposes *always* gives the salient set of possibilities to determine the directive principles that hold in a given factual context.

It could be that there are multiple sets of directive principles (specified relative to different sets of possibilities given by different sharpenings of 'feasible') all of which hold at once. Estlund has suggested that there could be many different sets of principles relative to different sets of facts applying at once. He is interested in the distinction between principles that are nonconcessive and principles that are concessive to people's moral non-compliance, and he notes that the nonconcessive standard requiring Professor Procrastinate to [accept and complete] applies at the same time as the concessive standard requiring him to refuse given that he will not complete. He suggests that it could be that 'there are infinitely many concession-relative standards, one for each set of moral shortfalls that are being taken for granted'.[239] It could be that the same is true more generally, that there are many different sets of principles relative to different sets of facts being held fixed (i.e., to different FCs). Perhaps it can be true simultaneously that you ought to φ (relative to FC f_1) and that you ought not to φ (relative to FC f_2). However, if this is the case, then another similar question arises about which is the salient set of possibilities (the salient FC) relative to which the directive principles telling us what we ought to do *all things considered* are specified. Given the multiplicity of different possible FCs and the lack of any single one that is obviously privileged, no simple, general answer to this question suggests itself. It may well be that there *is* no simple, general answer. It could be that the answer varies depending on the specific practical decision that we are facing. One upshot, then, of my account of feasibility is that there is significant uncertainty about what the feasibility constraint is relative to which the all-things-considered directive principles that apply to us are specified (given a factual context).

---

[239] Estlund (2017a) 44

This leads to the first part of the conclusion of this thesis. I set out to investigate whether it is possible to justify the common assumption that state legitimacy (the general and exclusive permission to enforce that states claim and are often taken to have) is on safe ground for at least some actual states and can be taken for granted (unlike the obligation to obey, which is quite widely thought not to exist). The combination of the argument of the first part of the thesis and the account of feasibility seems to have the result that it *cannot* simply be taken for granted. It cannot simply be *assumed* that this sort of general and exclusive permission to enforce is something that states (at least of the best existing kind) have. The first part of the thesis argued that plausible arguments for state legitimacy (or the legitimacy of states of a certain kind) depend on feasibility premises. These arguments tend to require a premise along the lines of the claim that the general and exclusive enforcement of the state is the only feasible way of meeting some moral requirement, or that the best feasible worlds involve the general and exclusive enforcement of the state. Given the above-mentioned uncertainty about what the feasibility constraint is relative to which the all-things-considered principles about moral requirement and permissibility are specified, it follows that in order to justify taking state legitimacy for granted, more work will have to be done to show that, on some feasibility constraint that *is* the salient one for the relevant all-things-considered directive principles (and so one on which something like the above principle P2 holds all things considered), there really is (for example) no other feasible way of meeting some moral requirement than through the general and exclusive enforcement of *this particular* state. (I argued also that the necessary feasibility premises do not obviously require an outlandishly unrealistic understanding of 'feasibility' to be false.)

However, given this uncertainty about the correct all-things-considered directive principles, we might ask a different question. Instead of asking what the correct all-things-considered principles are, we might ask what sets of principles (i.e., those relative to which sets of possibilities) are worth knowing about? Which sets of principles is it useful seeking to discover? It need not be the case that the sets of principles that it is worthwhile identifying are just those that hold for us.[240] It could be that there is a multiplicity of different sets of principles relative to different FCs that hold for us, but that there is no value in knowing about many of these. Conversely, it could be that the only set of principles that can be said genuinely to hold for us are the all-things-considered principles which are relative

---

[240] Cf. Swift (2008) 368.

to some determinate FC, but that nevertheless it is worthwhile identifying the sets of principles that would be true if *other* FCs were salient.

Above I said that one might have reason to think that the feasibility premises necessary for an argument for some state's legitimacy to go through can be taken for granted if one thinks that the sharpening of feasibility that constrains the truth of moral principles is some quite realistic one. However, I have just argued that given my account of feasibility, we have no immediate reason to suppose that this must be the case. One might, though, think that only principles specified relative to quite *realistic* feasibility constraints can be *of interest* or *of practical use*. There is no point asking about principles relative to FCs that are too unrealistic. In the next section, then, I will argue that this is wrong: it can be worthwhile and of practical use identifying principles constrained only by relatively unrealistic feasibility constraints.

3. Unrealistic theory: pro tanto principles

It is a common complaint against unrealistic moral or political theory (that is, theory unconstrained by realistic feasibility constraints) or against 'ideal theory' that it fails to be *action-guiding* and so is not interesting or useful.[241] Let us suppose, if only for the sake of argument, that moral theory that is worth doing must be action-guiding, that is, must contribute to the provision of guidance for our all-things-considered choice of real-world actions. Even so, I think that unrealistic theory *can* be action-guiding, despite its not being realistically constrained. I will, in this section and the next, argue that two different sorts of unrealistic moral theory can in general be useful for action guidance. I will then, in section 5, say something about why unrealistic theory specifically about the permissibility of state enforcement can be interesting and useful.

The first way in which I think unrealistic theory can be useful for action guidance relies on a somewhat speculative picture of moral principles (which I will suggest has some plausibility but I will not give a full defence of) and may be rejected if

---

[241] This is the objection discussed by Valentini (2009). The criticisms of ideal theory presented by, e.g., Sen (2006) and (2009), Schmidtz (2011), Goodin (1995) and (2012a), Farrelly (2007), Huemer (2016), Phillips (1995), Miller (2008) and Wiens (2015a) and (2015b) are all variants of this objection (or include variants of this objection). (This objection also forms part of the 'realist' challenge to ideal theory; see, for instance, Galston (2010) 395-6.) There are various persuasive defences of ideal theory as action-guiding, such as Valentini (2009), (2011) and (2012b), Swift (2008), Stemplowska (2008) and Mason (2004) and (2016). Here I defend *unrealistic* theory more specifically, and extend this defence to include theory about permissibility and requirement generally, not just in the domain of *justice*. Not all of the objections in the above works will be addressed here for lack of space; I focus on what I consider to be the most powerful such objection.

that picture is rejected, but in the next section I will offer another way in which unrealistic theory can be useful for action guidance, which is not dependent on this picture.

An obvious defence of unrealistic theory is that we need it to act as a target for which to aim, or as a benchmark by which to measure the moral acceptability or goodness of various realistic options. When we know what we would be required to do when a wider range of the current facts can vary, we will have an idea of what sorts of outcomes we should direct our action towards achieving. Though our action will be constrained here and now (in the short term) by low (realistic) FCs which may rule out the target, we can choose actions within these low FCs with an awareness of what would be better and thus what we should strive towards. Theorising at relatively high (unrealistic) FCs tells us what is 'ideal', or relatively ideal. This, it might be thought, gives us action guidance for what we should do, even though the theory is arrived at assuming that certain facts can change that it is unrealistic to think could change, because we should just attempt to get as close as possible to doing what we should do ideally. If, say, what we should do when we allow human motivations and so on to vary is achieve perfect equality of welfare, then when human motivations and so on are not variable and we cannot achieve such equality, what we ought to do, the thought goes, is get as close as possible to this ideal. Furthermore, the principles that represent the 'ideal' might be able to serve as a benchmark: how far what we are actually doing is from this ideal shows us how good or morally acceptable it is. Rawls's defence of ideal theory is an argument of this sort. 'Until the ideal is identified', he says, 'nonideal theory lacks an objective, an aim, by reference to which its queries can be answered'.[242] Other variations on this argument are present in various other contributions to this debate.[243]

However, it has been noticed that there is a result in economic theory that applies generally to the relation between optimal outcomes and 'second best' outcomes, which seems to pose a problem for this argument. In 1956 Richard Lipsey and Kelvin Lancaster proved a theorem in economics that they called 'the General Theory of Second Best'.[244] The idea is that if a Pareto optimal outcome consists in the fulfilment of a number of 'Paretian conditions', then

> given that one of the Paretian optimum conditions cannot be fulfilled, then an
> optimum situation can be achieved only by departing from all the other Paretian

---

[242] Rawls (1999b) 90. See also Simmons (2010).
[243] For instance, Stemplowska (2008) and Gilabert (2012).
[244] Lipsey and Lancaster (1956)

> conditions ... Specifically, it is *not* true that a situation in which more, but not all, of the optimum conditions are fulfilled is necessarily, or is even likely to be, superior to a situation in which fewer are fulfilled.[245]

Several philosophers have noted that the theory of second best (TSB) has implications for political philosophy.[246] Robert Goodin notes that the very strong conclusion Lipsey and Lancaster arrive at (that an optimum situation can be achieved *only* by departing from *all* the other Paretian conditions) stems from certain assumptions they made. However, the second point they make in the quote above holds more generally, he suggests. The second-best state of affairs is not *necessarily* identical to the first in any respect.[247] The idea, then, is that if our ideal or unrealistic theory calls for something (a state of affairs, a set of institutions, a set of principles) that has several features, but the constraints of the world prevent all of these features from being achieved *together*, the second-best alternative (or the best given the constraints) is not necessarily going to be just the one that is closest to the first-best in the greatest number of these features. If one feature is absent, it is not necessarily best to realise the others. The alternative that changes the least features of the ideal may actually be worse by the lights that led us to choose the first-best than one that changes more. Thus, if we take the best principles given some reasonably unrealistic FC as an ideal, it will not necessarily be the case that the best principles or outcomes given a more realistic FC will resemble them. Thus, if an unrealistic FC allows facts to change that it is very implausible to think could change, the thought goes, a theory constrained only by this unrealistic FC will not be much use as a target or standard for guiding action, since there is no guarantee that what we should do given a more realistic FC will resemble the target. Thus, it is not obvious how knowing what the target is could help us in judging principles or policies in the real world.

Goodin gives the analogy of a choice of car. Suppose my ideal car would have three features: it would be silver, new and a Rolls Royce. Suppose now that such a car is unavailable, but two others are. One is a week-old black Jaguar and the other is a new, silver Toyota. The latter has two of the three features of my ideal car, while the former has none. However, it is likely that I would in fact prefer the

---

[245] Ibid. 11-12
[246] The TSB is discussed explicitly by Estlund (2017b), Goodin (2012a) and (1995), Mason (2004), Swift (2008) and Wiens (2016). Similar challenges for 'ideal' political philosophy are discussed in Brennan and Pettit (2005), Phillips (1985) and Wiens (2015a).
[247] Goodin (2012a) 157. Estlund (2017b) notes that there is a 'fallacy of approximation' that is broader in scope than the TSB.

Jaguar and not the Toyota.[248] Since unrealistic theory involves assuming away certain constraints that will constrain our actions in the real world, these latter are bound to be constrained by things that did not constrain the unrealistic theory. Thus, it may often not be feasible, given more realistic FCs, to do what all of the principles relative to unrealistic FCs would require us to do. The TSB tells us that when one of the features of a recommendation is not present (or one of the principles is not satisfied), it is not necessarily better for the other features to be realised (or the other principles satisfied). This appears to show that unrealistic theory will not necessarily be a useful guide to what the best alternative is given more realistic constraints. This, then, suggests that identifying a target may in fact not be action guiding at all. (That is, when the target itself cannot be realised, it may not provide guidance for our all-things-considered choice of action.)

I think that the problem of second best is a problem for certain sorts of unrealistic theory that needs to be taken very seriously. However, I think that *only certain sorts* of unrealistic theory are vulnerable to this problem at all.[249] Once we distinguish different types of moral and political theory that might be done unconstrained by realistic feasibility constraints, we see that the problem of second best is not *generally* a problem for unrealistic moral and political theory. Several writers have made versions of this point, and in this section, I develop, and extend to the domain of deontological principles, a defence of a kind of unrealistic theory similar to those given by, for instance, Adam Swift, Zofia Stemplowska and Andrew Mason.[250]

If we conceive of unrealistic theory as doing something similar to choosing one's ideal car, then we come up against the problem of second best. My ideal car is a (possibly non-existent) car that is the best car I could imagine. The task of identifying one's ideal car is essentially a task of *design*. One designs a car, exactly the way one would like it to be. If we think of unrealistic theory as just like this, except for society instead of for a car, then the TSB poses a problem exactly analogously to how it does in the case of the car. Call this unrealistic *institutional design*.[251] In unrealistic institutional design we specify, as exactly as possible, how society should be, given an unrealistic feasibility constraint (i.e., given that

---

[248] Goodin (1995) 53 and Goodin (2012a) 157

[249] Wiens (2016) notices that certain interpretations of the application of the TSB do not pose a threat to the view of ideal theory as a target.

[250] Swift (2008); Stemplowska (2008); Mason (2004).

[251] This term is potentially misleading since, as well as institutions, this type of theory could involve designing policies or specific actions. For this reason, Robeyns (2008) prefers the term 'action design'. I keep the term 'institutional design' simply because it seems to be more widely used.

many/most/all options are feasible), designing the institutions that would create the best society possible. This ideal society then, like the ideal car, will have a number of attributes. Common sense might say that in order to achieve the best society we can within certain non-ideal constraints, we should create a society that instantiates as many as possible of the attributes of the ideal society to as great an extent as possible. However, what David Wiens calls an 'anti-approximation warning', which follows from the TSB, seems to hold in this case.[252] We cannot assume at all that the best thing to do, given the unachievability of the ideal society, is to create a society that instantiates more rather than less of the attributes of that ideal society. If this is the case, then this sort of ideal theory seems to give us very little guidance as to what to do in a world in which the fully ideal society is not achievable.

However, this is not the only sort of moral or political theory that might be worth pursuing unconstrained by realistic feasibility constraints. Hamlin and Stemplowska distinguish what they call 'theory of ideals' from ideal and non-ideal theory.[253] The purpose of theory of ideals, according to them, is to 'identify, elucidate and clarify the nature of an ideal or ideals'. This includes both an element 'devoted to the identification and explication of individual ideals or principles' and another 'devoted to the issues arising from the multiplicity of ideals or principles (issues of commensurability, priority, trade-off, etc.)'.[254] They describe this form of theory asking us to imagine a graph plotting the realization of two (or more) values or principles against each other. The task of the theory of ideals then involves both specifying the axes (that is, identifying what the values and principles *are*) and then identifying the shape and position of the indifference curves (that is, identifying between which bundles of realization of different values and principles we are indifferent). Thus, on a simple model with only two values, say equality and security, this would involve analysing what these values are or what they involve and deciding how they should be balanced when there is a limit to how much we can achieve of each. The task of theory of ideals, then, is roughly to identify the ideals that we aim to realise when designing institutions (as well as how they should interact with each other).

---

[252] Wiens (143). Wiens notes that the TSB need not pose a problem for ideal theory in general; it does not clearly show that we should not necessarily aim to satisfy ideal principles in the actual world. But he argues that it *does* show that, since we cannot assume we ought to approximate an ideal, we cannot assume that 'a political ideal presents an appropriate target for real-world reform efforts' (143). This complaint, though, as we will see, is only relevant when the 'ideal' is conceived of in certain ways.
[253] Hamlin and Stemplowska (2012)
[254] Ibid. 53

Hamlin and Stemplowska present the theory of ideals as concerned only with that part of morality to which optimising considerations apply. I think, though, that there is a part of moral theory that includes something like theory of ideals alongside a deontic component. Theory of ideals is not subject to feasibility constraints in the same way that all-things-considered directive principles, or institutional design, are, since its aim is to identify what our ideals and values *are*, what they consist in, rather than what we ought to do. I think, though, that the task of identifying ideals and values is part of a wider form of theory which includes the task of identifying a certain kind of deontic principle that is similarly not subject to feasibility constraints (call these the pro tanto requirements). Let us call 'theory of principles' the more general task that involves identifying not just ideals, but considerations that have moral weight more generally, which, as I will now suggest, can include pro tanto requirements and permissions.

As we saw above, there is a sense in which Professor Procrastinate *ought* to accept the assignment even though he will not complete it: he ought to [accept and complete] it. This seems to be a *pro tanto* ought. The fact that he ought to [accept and complete] does not figure in what he ought to do all things considered, but it nevertheless remains true that it is something that he ought to do. Wiens explains this by suggesting that such nonconcessive requirements are specified relative to a set of possible worlds that includes both worlds at which the requirement is not met and worlds at which it is. A requirement specified relative to a set of worlds holds at all of the worlds in this set, he suggests, even those at which it will not be met. It follows that requirements specified relative to the set of all worlds or other expansive sets of worlds that include the actual world are requirements that hold for us. It seems like these will just apply to us *pro tanto*. (There is some sense in which such pro tanto requirements are not really *requirements*: in some sense we are not *required* to follow them; it may be the case that you *ought not* do what you ought *pro tanto* to do. Nevertheless, regardless of the felicitousness of the term 'requirement', there seems to remain some sense in which these *are* moral 'oughts' that apply to us. I merely suggest that the domain of moral principles that have application is not exhausted by those we ought all-things-considered to follow; how exactly we should understand these pro tanto 'oughts' is a difficult question.)

This picture is supported by consideration of ordinary pro tanto 'oughts'. Suppose, for example, that you have made a promise to a friend to meet at 3 but on your way, you come across a stranger in danger of death who urgently needs your help to get to hospital; if you take this person to hospital you will be late to meet your friend. In a sense, it remains the case that you *ought* to meet your friend

at 3: this is a pro tanto 'ought'. But the pro tanto 'ought' is defeated by the stranger's need for your help. Relative to a set of possibilities that includes a world at which you both take the stranger to hospital *and* meet your friend at 3 you ought to do both. Given that a feasibility constraint that determines such a set of worlds also includes the actual world in the set, this pro tanto 'ought' holds at the actual world (as we ordinarily judge). Relative to a set of possibilities that does not include a world where you do both, though, it is not the case that you ought to meet your friend at 3. This gives us the all-things-considered claim that you ought *not* to meet your friend at 3.

In the same way, then, in which Procrastinate ought to accept and complete his assignment, we ought to do that which is required of us given various unrealistic feasibility constraints. That is, we ought to do these things *pro tanto*. The most expansive feasibility constraint holds fixed no facts, so on this sharpening anything is feasible for me that I bring about in the right way at *some* possible world identical to the actual world up until the present moment. This set of worlds obviously includes the actual world, and so principles specified relative to it hold for us (pro tanto). There is thus a domain of pro tanto requirements and permissions that are specified relative to the most expansive feasibility constraint, and to various other unrealistic feasibility constraints. I think that the task of identifying these pro tanto principles is a part of the task of identifying the moral considerations that should have weight in our actual all-things-considered decisions about what to do in just the same way as the identification of ideals and values is a part of this task. When we decide what to do, numerous considerations enter into the decision including both ideals whose realisation we seek to maximise (insofar as allowed by the countervailing importance of maximising other values and of respecting pro tanto requirements) and pro tanto requirements and permissions that we seek to respect (insofar as allowed by other countervailing requirements and by the importance of certain values).

Now, this task, the identification of ideals and pro tanto requirements, which I am calling 'theory of principles', is not constrained by feasibility considerations.[255] It is thus in some sense 'unrealistic'. It might seem, then, that it is subject to something like the problem of second best. There is no reason to suppose, the thought would go, that the pro tanto principles bear any resemblance to the all-things-considered principles (or that the ideals as specified by theory of principles

---

[255] Hamlin and Stemplowska (2012) make this point. Gheaus (2013) also argues that *justice* (the identification of which will be part of theory of ideals, and consequently, theory of principles) is not constrained by feasibility. Mason (2004) makes a similar point.

bear any resemblance to what maximisation of the ideals would require of us in the actual world). It would thus seem like the task of theory of principles is of no action-guiding use. However, this is to miss out a crucial part of the theory of principles. As noted above, Hamlin and Stemplowska describe 'theory of ideals' as involving two components, one of which is 'devoted to the issues arising from the multiplicity of ideals or principles'. For them, this involves something like identifying indifference curves, i.e., identifying between what bundles of different values we should be indifferent. When we generalise to theory of principles, the task cannot easily be described as the identification of indifference principles, but there is still an analogous task concerning the interactions between different principles and ideals.

Part of the task of theory of principles is to identify balancing principles that tell us how to balance the various pro tanto requirements and ideals against each other. As Hamlin and Stemplowska describe theory of ideals, it will tell us what we ought to do in a concrete factual situation by telling us simply to realise the highest indifference curve possible within the relevant feasibility constraint. Something similar should be true for theory of principles: balancing principles should tell us how to weigh ideals and pro tanto requirements against each other given a feasibility constraint.[256] Thus, in principle, if this task were complete, theory of principles should tell us how to balance these various considerations for any given factual context and for any given feasibility constraint. This is not to deny that the task is so complex that it is unlikely ever to be complete. That it is a complex and difficult task (and unlikely ever to be complete) is no reason to suppose that insofar as it *can* be done it is not worth doing. Nor is it even to claim that there *are* complete balancing principles.[257] It could well be that there are some intractable conflicts or incommensurabilities between values. This would mean that this task could not be complete in the relevant sense; there would be no set of pro tanto requirements, ideals and balancing principles that always gives a determinate answer to the question what to do. Nevertheless, identifying the pro tanto requirements, ideals and balancing principles that it *is* possible to identify would tell us what moral considerations apply in which situations and how to balance them insofar as they can be balanced.

It should thus be clear, I think, that the problem of second best does not prevent this form of unrealistic theory from being useful for action guidance. The anti-

---

[256] Cf. Mason (2004) 265
[257] Galston (2010) appears to offer this as an objection to 'ideal theory' (407), but for the reasons adduced here, I cannot see that it poses a general problem at all if ideal theory is understood as theory of principles.

approximation warning just does not apply here: the point of this sort of unrealistic theory is not to identify some model that we ought to approximate as much as possible. It is rather to identify the values, ideals and pro tanto requirements that are of moral importance.[258] To know what we ought to do in a given factual situation, it will presumably be useful to know what considerations have moral weight: what has value worth realising and what pro tanto requirements we are under. It is thus not in general the case that unrealistic theory cannot be worthwhile or useful. I will return below to what this means for state legitimacy more specifically.

4. Unrealistic theory: all-things-considered principles

The argument in the previous section relied on a somewhat speculative picture of moral principles in order to argue that a certain kind of unrealistic theory is useful for action guidance. I have not offered a full argument for this picture, and you may not want to accept it. I will now, though, argue that another kind of unrealistic moral and political theory can be worthwhile, and this argument will not depend on the above picture of moral principles. In the above I distinguished theory of principles from institutional design and argued that the former is not subject to the problem of second best. The latter, though, along with the identification of all-things-considered directive principles (from now on, for the sake of brevity, I will use 'institutional design' to refer to the conjunction of these two things), telling us concretely what to do here and now, does seem vulnerable to the problem of second best. In carrying out institutional design, the results we get will depend on what FC we take as a constraint. If we vary the constraint, different institutions or actions will come out as morally best or morally required. Of course, not all of these possible sets of institutions and actions can be those that we actually ought all-things-considered to implement. If it were not for the theory of second best, it could appear that what we ought to do is simply carry out institutional design constrained only by a highly unrealistic FC and then attempt to approximate the resulting institutions and actions as best we can. However, the theory of second best tells us that if we do not succeed in fully implementing the unrealistic institutional design, the best approximation we do manage to achieve may actually be worse than other available alternatives. It might seem, then, that when we attempt to identify all-things-considered principles for action and institutional design we should always do so within a

---

[258] Goodin seems, in fact, to realise that the problem of second best only has application for something like institutional design and not for theory of principles (see Goodin (1995) n. 45).

quite realistic feasibility constraint. However, I think that despite the theory of second best, relatively unrealistic institutional design or all-things-considered principles *can* be useful for action guidance. Identifying what we ought to do (or what we *would* be required to do) relative to unrealistic FCs can be helpful in identifying what we ought to do relative to more realistic ones.

As we saw above, it is a consequence of my account of feasibility that there is no single sharpening of feasibility that is obviously the one relative to which all-things-considered principles are to be identified. It is an open question *exactly* how realistic we ought to be when deciding what to do, or how to structure our institutions. In addition, it is not apparent how one would go about attempting to answer this question. Indeed, I suspect that there is no single general answer. But, if we cannot easily say which FC constrains the answer to any given all-things-considered practical question, can we say anything about which FCs it can be *useful* to take as constraints on institutional design? If, contra the theory of second best, approximating to an unrealistic institutional design were always better than alternatives, it would clearly be useful to carry out unrealistic institutional design even though its outputs are not themselves the actions and institutions we ought all-things-considered to implement. Even though this is not the case, it is clearly possible for institutional design to be *useful* even if its outputs are not themselves what we ought to implement.

Given that it is not obvious which FC constrains all-things-considered principles, we seem to face a dilemma when deciding what to do, or what institutions to create, all-things-considered. On the one hand, we could just focus on doing what we ought to do given a relatively realistic FC. Since we want to guide action in the short-term, where relatively little is changeable, this might seem to be what we should do. On the other hand, however, we want to be more demanding than this. We cannot be exclusively restricted to short-term (realistic) FCs, since this is to ignore what Gilabert calls 'dynamic duties'.[259] We may have a dynamic duty to do x if it is not now possible to do x, but it would be possible to do x if we did y and we *can* now do y. If x is something we should do and it is worth the cost of doing y, then we should do y. Because of this, Gilabert thinks, political philosophy should adopt a 'transitional standpoint', which 'focuses on the identification of dynamic trajectories of political action, which set into motion a sequence of political reforms passing through successive thresholds of feasibility'.[260] Thus, we

---

[259] Gilabert (2011, 59ff) and (2012, 47ff)
[260] Gilabert (2008) 411. Stemplowska (2008) and Simmons (2010) also underline the importance of transition.

might prefer to design institutions on a more unrealistic FC in order to avoid excessive conservatism and to get the best institutions allowing for some greater changes to be made.[261] However, the theory of second best shows us that focusing on a more distant goal in this way may, if we fail to achieve that goal, lead us to do things that are not in fact the best things we could have done. Thus, on the one hand, pursuing what is desirable given realistic FCs may prevent us from doing better relative to less restrictive FCs. On the other hand, focusing on more unrealistic FCs might lead us to end up with worse outcomes than we could have had.

I think that the obvious way around this dilemma is to be pluralistic in our theorising.[262] The best way to avoid unnecessary conservatism and excessive optimism is to identify what the all-things-considered principles for action (or institutional design) would be relative to a variety of different more and less realistic FCs. If there were no constraints on the time we had to devote to theorising, we would identify what these principles would be relative to each and every FC. We would then attempt to assess the plausibility of our achieving what is required of us on these various FCs. With this knowledge, we can weigh up the potential short-term losses that would come from aiming for something we cannot in fact achieve (weighted according to how plausible it is that we will meet the unrealistic requirement) with the long-term gains that would come from achieving the unrealistic target (similarly weighted). We obviously do not have the capacity to carry out this sort of theory relative to *every single* FC, since the number of FCs is, if not infinite, at least massive. But in general, having an idea of what we ought to do given a variety of FCs of varying degrees of realisticness will help us to weigh up short-term loss against long-term gain and come to some sort of intuitive judgment about what we actually ought to *do*, which feasibility constraint we ought to take as operative for a given practical decision.[263] Sometimes we ought not to do what would be required on a realistic FC because of what not doing so might make possible in the long term, while equally sometimes we ought to do something that makes impossible the best outcome in the long term (what would be required given an unrealistic FC) because of the

---

[261] Valentini notes that the more realistic the FC, 'the more [the theory] will appear to offer an uncritical defence of the *status quo*' (2012b) 659.

[262] Valentini (2012b) makes the related point that the appropriate level of 'realisticness' will depend on the purpose of one's theorising.

[263] Carens (1996) and Valentini (2012b) have both argued that theory at various points along an 'ideal/non-ideal' scale can be worthwhile. As Carens says, 'there is no single correct starting point for reflection, no single correct set of presuppositions about what is possible ... The assumptions we adopt should depend in part on the purposes of our inquiry' (169).

badness or impermissibility of what would be required in the short term to realise the long-term best outcome.

Thus, it seems that identifying principles for action or institutional design relative to reasonably unrealistic feasibility constraints *can* be useful and even important. Doing so is important to avoid undue conservatism, and a helpful part of determining exactly *how* conservative we ought to be. In addition, if we discover that something is required or permitted given a more expansive FC that is different to what is required or permitted given a restrictive, realistic FC, this shows us that it is worth considering what we can do to make it possible to realise the former principle. In other words, we learn that it is worth investigating what can be done to push us onto a more unrealistic feasibility constraint (so to speak). More unrealistic feasibility constraints give us more options; that is, they make more things come out as feasible. Thus, realising the principles that hold relative to more unrealistic FCs is in some sense better, more ideal. Thus, if something is required on a more unrealistic FC, we have a reason to see if something can be done to meet that requirement.

5. Unrealistic anarchism

Where has all this got us? I have just argued that it can be worthwhile doing moral or political theory constrained only by unrealistic feasibility constraints: either to identify the ideals and pro tanto requirements that have moral weight or simply to avoid undue conservatism when deciding what concretely to do. What does this mean for the legitimacy of existing states? My thesis has argued that plausible arguments for the legitimacy of existing states require feasibility premises. As we saw above, this has the consequence, given my account of feasibility, that we cannot simply take state legitimacy for granted: some argument is needed to show that the sharpening of 'feasibility' relative to which the relevant all-things-considered principles are specified is one on which the feasibility premises needed to establish state legitimacy are true. However, I said, since we are left with uncertainty about which set of principles (those specified relative to which feasibility constraint) are those that tell us what we ought, all things considered, to do, we could instead ask which sets of principles it is *worthwhile* identifying. The conclusion in this chapter was that it is not only realistic principles (those constrained by restrictive feasibility constraints) whose identification can be useful for action guidance.

It was part of the claim of the first part of my thesis that successful arguments for state legitimacy depend on feasibility premises that will not be true on *all* possible

sharpenings of feasibility. Thus, there is a certain level of unrealisticness at which it ceases to be the case that legitimacy could be established for existing states. I have argued that, *in general*, unrealistic theory can be useful for action guidance. But what *practically* do we learn when we learn that the principles specified relative to some (at least somewhat) unrealistic FCs do not (or would not) give states the sort of general and exclusive permission to enforce that they claim?

As we saw above, there are two kinds of unrealistic theory, and the conclusion that arguments for state legitimacy depend on feasibility premises is relevant to both kinds of theory. First, I take it that when we learn that this sort of permission to enforce cannot be established for any state a priori, we learn that the general and exclusive permission to enforce (for states that meet certain conditions) is not a part of the domain of theory of principles (it is not a part of the domain of pro tanto requirements specified relative to the most expansive feasibility constraints).[264] What use is this knowledge? There are several answers to this question. The first is not straightforwardly a way in which this knowledge is useful for action guidance, but there may be an indirect way in which the use it alludes to is relevant to decisions about what to do. The knowledge that the more fundamental moral considerations that apply to us do not include a general and exclusive permission to enforce for any existing states is relevant to how we ought to think about and treat our states' enforcement. We should not think of our states' general and exclusive enforcement as something *good* in itself, and to be welcomed (as, for example, Pettit wants to say it is).[265] It may be that in some sense it is good and to be welcomed, as something that makes possible something good, but at best, contra Pettit, it is a necessary evil, a remedy for unfortunate circumstances. We should think of (and treat) the state's exclusive enforcement as only justified insofar as there is some other value that cannot be realised without it, or some other moral requirement that cannot be met without it.

That the state's general and exclusive permission to enforce does not show up at the level of theory of principles shows us that this is something where we should do better *if we can*. If we were not subject to certain feasibility constraints, we would not be permitted to have a state with general and exclusive enforcement

---

[264] It does not quite follow from the fact that legitimacy cannot be established a priori that such a permission cannot form a part of the domain of theory of principles, since it could be that some a posteriori facts that are not to do with feasibility would suffice to establish it, and then it could still form a part of the domain of theory of principles in worlds at which the relevant a posteriori facts hold. However, my arguments seemed to show not only that legitimacy cannot be established a priori, but that the kind of a posteriori premises necessary to establish it include feasibility premises.

[265] Pettit (2012) 181

power. Thus, *if it were possible* to remove these feasibility constraints, that would be something that we should do. From this, two things follow. First, it follows that this is a matter where we fall short of the ideal. Unfortunate circumstances make it infeasible to do better, perhaps, but what we learn is precisely that this is *unfortunate.*

Robert Jubb argues forcefully that one reason full-compliance theory is worth doing is because it shows us when and to what extent partial-compliance theory is *tragic:*

> The reason that we should do ideal theory in the sense of full-compliance theory is because, without doing it, we will often not be able to understand if and to what extent non-ideal theory is tragic. In not understanding that tragedy, we will tend to make mistakes about what is actually desirable in the circumstances of tragedy. One of the important normative features of those circumstances is their tragedy, and being unable to grasp that feature will tend to generate mistakes about how to respond to it. Unless we know how people are being mistreated, then we are likely to continue to mistreat them and fail to take the steps we are obliged to in order to rectify their mistreatment.[266]

I think this point generalises to unrealistic theory more widely. When we learn that our state's enforcement violates a pro tanto moral requirement, we learn that, if it *is* permissible all things considered, this is a tragic fact. (When we break a pro tanto moral requirement because it is impossible to meet it alongside some other weightier requirement, there may be no *all-things-considered* question about whether it is the right thing to do, but there is still a moral loss.) It then becomes important to think about *why*, if at all, it is permitted, and exactly *when* it is permitted. *What*, if any, are the feasibility constraints that make generalised state enforcement permissible, and when do they apply?

The second thing that follows from the fact that states' general and exclusive enforcement is not permitted by theory of principles is that if there *is* something we can do to make the ideal possible, we should. It is thus a consequence of this 'unrealistic anarchism' that it is worth devoting some thought to what might be done to make it possible to meet the pro tanto requirements that generate a complaint against state enforcement concurrently with the pro tanto requirements that in certain circumstances make state enforcement permissible (in other words, what might be done to achieve a morally acceptable anarchy). If we ought to avoid generalised state enforcement where compatible with meeting other pro tanto requirements, then it is worth exploring

---

[266] Jubb (2012) 238-9

whether there *is* a feasible way of doing so.[267] It could be that there are certain domains in which we could more easily do without the state's generalised enforcement. We ought to avoid violating the pro tanto requirement ordinarily violated by state enforcement in as many domains as possible without jeopardising other pro tanto requirements.

One further way in which the knowledge that a state's legitimacy is not a part of the domain of theory of principles might be useful is as a form of inspiration. Knowledge of the more fundamental moral requirements and ideals can have a role *inspiring* our short-term action. This knowledge might inspire us, for instance, to do our best to meet other moral requirements or realise other values *without* state enforcement. Further, it might inspire us to treat our states with suspicion and to be vigilant as to whether their exclusive and general enforcement really is necessary to meet some moral requirement and to resist, for example, if the state *fails* to meet the moral requirement that might make its enforcement permissible. We may be inspired to treat the state's exclusive enforcement as something only to be accepted *insofar as necessary* for some purpose (the realisation of some ideal or the meeting of some moral requirement).

There was a second form of unrealistic theory that I discussed above: unrealistic institutional design. I argued that such theory can be practically useful as a contribution to the process of determining what we ought, all things considered, to do, achieving a balance between undue conservatism and excessive optimism. When we identify that states lack an exclusive and general permission to enforce given a number of more expansive feasibility constraints (and so, that given those constraints, we ought not to have a state with generalised enforcement power), we learn something that we can factor into deciding what to do in the way described above. Thus, even if we reject the picture of theory of principles suggested above, this conclusion has practical implications for our all-things-

---

[267] There are, of course, constraints on our time and, as Robert Goodin (2012b) argues, sometimes it is necessary to take certain things as settled. To be able to make plans and have goals, and to render our behaviour predictable for others, we need to take certain matters as fixed. Sometimes we need to just accept certain moral failings. However, as Goodin agrees, we should not accept *everything* as settled: we need to *strive* as well. Indeed, it seems to me that we ought to strive to do better morally as far as is possible without radically threatening our ability to plan our lives. Goodin offers no recipes for determining when we ought to strive and when settle, but it could be that state enforcement is one of those things that we should take as fixed. There may be more urgent matters, and for certain purposes it is perfectly acceptable to take state enforcement as fixed. However, since it is of great significance to our lives, it does seem that it could be worth investigating whether there are ways to meet weighty moral requirements and realise weighty values without it. Even if it turns out not to be possible to do without it altogether, there may be certain domains in which we can.

considered action. These implications are broadly similar to those of the theory of principles conclusion above, although thinking in terms of institutional design puts the emphasis slightly differently. We learn that if there is something we can do (that is not too costly in other terms) to push us onto the more expansive feasibility constraints where a morally acceptable stateless world is feasible, then we should. This knowledge may again act as an inspiration.

When we learn that given a more expansive feasibility constraint, we ought to structure our societies without the state's exclusive and general enforcement, we learn that this is something that needs to be weighed up against what the costs of making this possible, or working towards a world where it is possible, would be. Knowing that in some sense a world without such enforcement is desirable should factor into our all-things-considered decisions. When it is possible to work towards such a world without excessive costs, we should do so. Like the theory of principles conclusion, this has the implication that it becomes a worthwhile task identifying *what* the feasibility constraints are that make such a world impossible (or impossible in a desirable way). What exactly would we have to do to get there (in a desirable way)? This may be a task more for social science than for normative philosophy, but the normative conclusion has the consequence that this social scientific task is an important one.

What *exactly* all of this implies for how we should act here and now will of course depend on *what* the moral requirements or ideals are that make state enforcement permissible (as the only feasible way of realising/meeting these ideals/requirements). It could be thought that the 'here and now' recommendations that will follow from my thesis will be in line with something like the Nozickian libertarian call for a minimal state that enforces property rights and basic rights against aggression, but does not engage in redistributive taxation or other state activities. What my thesis shows (if right) is that exclusive state enforcement is only justifiable insofar as *necessary* for meeting some important moral requirement, or for realising some important goal. It might seem to follow from this that we should *minimise* state enforcement. If we then were to argue that it is not feasible to achieve, say, rights-respect or rights-security without exclusive state enforcement, we then might be tempted to conclude that it is permissible all-things-considered for states to enforce libertarian basic rights, but nothing else. This is something like the Nozickian minimal state. It should be clear, though, that this *need not* be the case. We only get to this sort of libertarian minimal state from my conclusions if we make certain assumptions.

We get there if, but only if, we assume that people have strong property rights in external objects whose protection is of significant moral importance, sufficient to overcome the presumption against exclusive state enforcement. We would need to assume first that the sort of value or moral requirement that might necessitate state enforcement is the security of individuals' rights and not, say, distributive justice (of a patterned, or end-state, kind). If exclusive state enforcement is justified because it is the only feasible way of protecting rights, then it is plausible that states' enforcement will only be justified when it *is* necessary for rights protection. Further, it needs to be assumed that people can have strong property rights in external objects. Without this, even if we thought security was important and what justifies state enforcement, we might be able to get to the idea that we should seek a minimal state, but we would have no reason to think that this minimal state must be one that allows 'capitalist acts between consenting adults', one, in other words, that protects property rights and does not engage in redistributive taxation. We only have a reason to include 'capitalist acts' in the realm of the basic freedoms that minimal states have a justification for protecting if we think people have strong property rights.

A society's system of property rights is one aspect among others of the way social cooperation is structured that needs to be determined in some way collectively. The collective rules for social cooperation that a society settles on need not be *enforced*. But if it is in some sense *infeasible* for these rules to be observed without enforcement, this could be enough to warrant state enforcement of these rules if they are of sufficient value. However, if we think we ought to *minimise* state enforcement, the question still remains what we ought to minimise it *to*. The answer to this latter question will be a matter of which social rules we consider to be most morally important, which aspects of living together we consider worth the violation of the presumption against state enforcement. There is no ex ante reason to suppose that this must include strong property rights and the protection of 'capitalist acts'. We could instead think, for instance, that the minimal social rules we should enforce should include rules that ensure distributive justice, a basic economic minimum, or even full economic equality.[268]

---

[268] The 'left-libertarians' make the related point that full private property rights are not essential to a libertarian commitment to self-ownership. (See Vallentyne and Steiner (2000).) But, though the above is targeted at right-libertarianism which defends a minimal state limited to enforcing property rights and little else, my thesis need not lead us to left-libertarianism either, for similar reasons. If we do not share the left-libertarian belief in the importance of 'self-ownership rights', we need not take such rights as a part of the minimal domain of rules to which we ought to limit state enforcement. Self-ownership rights could be protected by collective rules for social cooperation, but they need not be: whether they should be is again a substantive moral question.

The collective social rules need not include strong property rights to external objects: they could instead, for instance, allow for exclusive property rights only insofar as consistent with some principle of egalitarian justice. We are only forced to go for a Nozickian minimal state if we think there are grounds for strong natural property rights.

However, the thought that my thesis calls for a minimal state is not entirely wrong. I think my thesis *does* suggest that we should, within some constraints, minimise state enforcement. But this leaves open what we should minimise it *to* (i.e., within *what* constraints we should minimise it). Since my thesis seems to show that state enforcement can only be justified insofar as it is the only feasible way (on some sharpening of 'feasible') of meeting some moral requirement or achieving some valuable goal, it does seem to follow that we should try and minimise state enforcement within constraints placed by weightier moral considerations. [269] What these constraints are, what we ought to minimise state enforcement to, will depend on your moral view, your theory of justice and so on. If we thought that the moral requirement that could justify state enforcement was not a requirement to protect individuals' rights, but rather some sort of requirement for egalitarian justice, we would get a very different picture of the sort of 'minimal state' we ought to strive for: one, presumably, that maintains egalitarian justice and little else, and that does so with a minimum of enforcement.

It might seem, then, that this makes my conclusion somewhat unhelpful, even redundant for practical purposes. This, though, would be wrong. My thesis does not offer a full set of instructions for what sort of state enforcement we ought, here and now, to support, or to resist. But it does, if I am right, offer a significant contribution to this task. What it claims to show is that we *should* minimise state enforcement, as far as we can without compromising weightier moral values or requirements. We ought to do what is possible to reduce state enforcement in line with our other moral commitments: we should not, contra Pettit and the Kantians, treat it as a good and welcome thing. According to the a priori view of state legitimacy, state enforcement (of certain kinds) will always be a good and morally essential thing; we ought not to seek to do without it, since otherwise there can be no right, or no freedom. It is a consequence of my view, however,

---

[269] The lesson of the theory of second best is that it does *not* follow from the fact that the ideal society would be one with no state enforcement that we ought to minimise state enforcement *tout court*. But it does follow from the fact that state enforcement violates a pro tanto requirement that we ought to avoid it whenever that requirement is not defeated by weightier considerations: in other words, we ought to minimise it within the constraints imposed by those considerations.

that the possibility (and the feasibility) of doing without the state, in as many domains as we can, is one worth exploring.

6. Conclusion

The two primary parts of my thesis, the argument that state legitimacy cannot be established a priori (and can only be established with the aid of feasibility considerations) and the multivocal account of feasibility, when combined lead to two headline conclusions. The first is that we cannot be confident that state legitimacy is on safe ground; it is not something that should be simply taken for granted. This is because we lack an argument that the sharpening of feasibility that constrains morality (or the relevant domain of morality) is one on which the necessary feasibility premises are true. The second is that, even if it turns out that what we are permitted to do all things considered is constrained by a reasonably realistic feasibility constraint, and so that some states *do* have a general and exclusive permission to enforce all things considered, it is still the case that these states would lack legitimacy given more unrealistic feasibility constraints and that their legitimacy is not a part of the domain of theory of principles. As I have argued in this final section, this is something that is worth knowing.

References

Anderson, Scott (2010) 'The Enforcement Approach to Coercion' *Journal of Ethics and Social Philosophy* 5: 1-31

-------------------- (2014) 'Coercion', *The Stanford Encyclopedia of Philosophy* (Spring 2014 Edition), Edward N. Zalta (ed.), URL = http://plato.stanford.edu/archives/spr2014/entries/coercion/

Anscombe, Elizabeth (1978) 'On the Source of the Authority of the State' *Ratio* 20: 1-28

Applbaum, Arthur Isak (2010) 'Legitimacy without the Duty to Obey' *Philosophy and Public Affairs* 38: 215-39

Arneson, Richard (2002) 'Democracy at the National Level' in Thomas Christiano (ed.) *Philosophy and Democracy: An Anthology* Oxford: Oxford University Press

Beitz, Charles (1979) *Political Theory and International Relations* Princeton, NJ: Princeton University Press

Berman, Paul (1972) (ed.) *Quotations from the Anarchists* New York City: Praeger

Brennan, Geoffrey (2013) 'Feasibility in Optimising Ethics' *Social Philosophy and Policy* 31: 314-29

Brennan, Geoffrey and Pettit, Philip (2005) 'The Feasibility Issue' in Frank Jackson and Michael Smith (eds.) *The Oxford Handbook of Contemporary Philosophy* Oxford: Oxford University Press

Brennan, Geoffrey and Sayre-McCord, Geoffrey (2016) 'Do Normative Facts Matter … To What is Feasible?' *Social Philosophy and Policy* 33: 434-56

Brennan, Geoffrey and Southwood, Nicholas (2007) 'Feasibility in Action and Attitude' in T. Rønnow-Rasmussen, B. Petersson, J. Josefsson & D. Egonsson (eds.) *Hommage à Wlodek: Philosophical Papers Dedicated to Wlodek Rabinowicz* www.fil.lu.se/hommageawlodek

Brighouse, Harry (2004) *Justice* Cambridge, UK: Polity Press

Buchanan, Allen (2002) 'Political Legitimacy and Democracy' *Ethics* 112: 689-719

------------------- (2004) *Justice, Legitimacy and Self-Determination: Moral Foundations for International Law* Oxford: Oxford University Press

Carens, Joseph H. (1996) 'Realistic and Idealistic Approaches to the Ethics of Migration' *International Migration Review* 30: 156-170

Clark, John P. (1978) 'What is Anarchism?' in J. Roland Pennock and John W. Chapman (eds.) *Nomos XIX: Anarchism* New York City: New York University Press

Cohen, G.A. (2009) *Why Not Socialism?* Princeton NJ: Princeton University Press

Copp, David (1999) 'The Idea of a Legitimate State' *Philosophy and Public Affairs* 28: 3-45

Cowen, Tyler (2007) 'The Importance of Defining the Feasible Set' *Economics and Philosophy* 23: 1-14

Durning, Patrick (2003) 'Political Legitimacy and the Duty to Obey the Law' *Canadian Journal of Philosophy* 33: 373-89

Edmundson, William A. (1998) *Three Anarchical Fallacies: An Essay on Political Authority* Cambridge, UK: Cambridge University Press

Elster, Jon (1985) *Making Sense of Marx* Cambridge and Paris: Cambridge University Press and Maison des Sciences de l'Homme

Estlund, David (2009) *Democratic Authority* Princeton, NJ: Princeton University Press

------------------- (2011) 'Human Nature and the Limits (If Any) of Political Philosophy' *Philosophy and Public Affairs* 39: 207-37

------------------- (2014) 'Utopophobia' *Philosophy and Public Affairs* 42: 113-34

------------------- (2016) 'What Is Circumstantial About Justice?' *Social Philosophy and Policy* 33: 292-311

------------------- (2017a) 'Prime Justice' in Michael Weber and Kevin Vallier (eds.) *Political Utopias: Contemporary Debates* Oxford: Oxford University Press

------------------- (2017b) 'Approximation, Deviation, and the Use of Political Ideals' paper draft presented November 8th 2017 to KJuris: King's Legal Philosophy Workshop, King's College London

-------------------- (2017c) 'Methodological moralism in political philosophy' *Critical Review of International Social and Political Philosophy* 20: 365-79

Farrelly, Colin (2007) 'Justice in Ideal Theory: A Refutation' *Political Studies* 55: 844-64

Foot, Philippa (2002) 'The Problem of Abortion and the Doctrine of Double Effect' in Philippa Foot *Virtues and Vices* Oxford: Oxford University Press

Galston, William (2010) 'Realism in Political Theory' *European Journal of Political Theory* 9: 385-411

Gheaus, Anca (2013) 'The Feasibility Constraint on the Concept of Justice' *The Philosophical Quarterly* 63: 445-64

Gilabert, Pablo (2008) 'Global Justice and Poverty Relief in Nonideal Circumstances' *Social Theory and Practice* 34: 411-38

---------------- (2009) 'The Feasibility of Basic Socioeconomic Human Rights: A Conceptual Exploration' *The Philosophical Quarterly* 59: 659-81

----------------- (2011) 'Debate: Feasibility and Socialism' *The Journal of Political Philosophy* 19: 52-63

----------------- (2012) 'Comparative Assessments of Justice, Political Feasibility, and Ideal Theory' *Ethical Theory and Moral Practice* 15: 39-56

---------------- (2017) 'Justice and Feasibility: A Dynamic Approach' in K. Vallier and M. Weber (eds.) *Political Utopias: Contemporary Debates* Oxford: Oxford University Press

Gilabert, Pablo and Lawford-Smith, Holly (2012) 'Political Feasibility: A Conceptual Exploration' *Political Studies* 60: 809-25

Goodin, Robert E. (1995) 'Political Ideals and Political Practice' *Journal of Political Science* 25: 37-56

----------------- (2012a) 'The Bioethics of Second-Best' in Joseph Millum and Ezekiel J. Emanuel (eds.) *Global Justice and Bioethics* New York City: Oxford University Press

----------------- (2012b) *On Settling* Princeton, NJ: Princeton University Press

Green, Leslie (1988) *The Authority of the State* Oxford: Oxford University Press

--------------------- (1996) 'Who Believes in Political Obligation?' in John T. Sanders and Jan Narveson (eds.) *For and Against the State: New Philosophical Readings* Lanham, MD: Rowman and Littlefield Publishers, Inc.

--------------------- (2004) 'Law and Obligations' in Jules L. Coleman, Kenneth Eimar Himma and Scott J. Shapiro (eds.) *The Oxford Handbook of Jurisprudence and Philosophy of Law* Oxford: Oxford University Press

Greenawalt, Kent (1999) 'Legitimate Authority and the Duty to Obey' in William A. Edmundson (ed.) *The Duty to Obey the Law: Selected Philosophical Readings* Lanham, MD: Rowman and Littlefield Publishers, Inc.

Greene, Amanda (2016) 'Consent and Political Legitimacy' in David Sobel, Peter Vallentyne and Steven Wall (eds.) *Oxford Studies in Political Philosophy, Volume 2* Oxford: Oxford University Press

Guillery, Daniel (forthcoming) 'Hobbes: A voluntarist about the permissibility of state enforcement?' *Ethics, Politics and Society*

Hamlin, Alan (2012) 'Political Feasibility' *e-IR* 29 August 2012, http://www.e-ir.info/2012/08/29/political-feasibility/ [accessed 23 February 2017]

Hamlin, Alan and Stemplowska, Zofia (2012) 'Theory, Ideal Theory and the Theory of Ideals' *Political Studies Review* 10: 48-62

Hampton, Jean (1986) *Hobbes and the Social Contract Tradition* Cambridge, UK: Cambridge University Press

Hawthorn, Geoffrey (1991) *Plausible Worlds* Cambridge: Cambridge University Press

Hobbes, Thomas (1996) *Leviathan* (ed. J.C.A. Gaskin) Oxford: Oxford University Press

Horton, John (2012) 'Political legitimacy, justice and consent' *Critical Review of International Social and Political Philosophy* 15: 129-48

Huemer, Michael (2016) 'Confessions of a Utopophobe' *Social Philosophy and Policy* 33: 214-34

Hume, David (1994) *Political Essays* ed. Knud Haakonssen. Cambridge, UK; Cambridge University Press

Ichikawa, Jonathan Jenkins and Steup, Matthias (2017) 'The Analysis of Knowledge' in Edward N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy* (Spring 2017 edition), URL = <https://plato.stanford.edu/archives/spr2017/entries/knowledge-analysis/>.

Jackson, Frank and Pargetter, Robert (1986) 'Oughts, Opinions, and Actualism' *Philosophical Review* 95: 233-55

Jensen, Mark (2009) 'The Limits of Practical Possibility' *Journal of Political Philosophy* 17: 168-84

Jubb, Robert (2012) 'Tragedies of non-ideal theory' *European Journal of Political Theory* 11: 229-46

Julius, A.J. (2017) 'Independent People' in Sari Kisilevsky and Martin J. Stone (eds.) *Freedom and Force: Essays on Kant's Legal Philosophy* Oxford: Hart Publishing

Kant, Immanuel (1999) 'The Metaphysics of Morals' in Mary J. Gregor (ed.) *Practical Philosophy (The Cambridge Edition of the Works of Immanuel Kant)* Cambridge, UK: Cambridge University Press

Kavka, Gregory S. (1996) 'Why Even Morally Perfect People Would Need Government' in John T. Sanders and Jan Narveson (eds.) *For and Against the State: New Philosophical Readings* Lanham, MD: Rowman and Littlefield Publishers, Inc.

Klosko, George (1992) *The Principle of Fairness and Political Obligation* Lanham, MD: Rowman and Littlefield

Kolodny, Niko (2014a) 'Rule Over None I: What Justifies Democracy?' *Philosophy and Public Affairs* 42: 195-229

-------------------- (2014b) 'Rule Over None II: Social Equality and the Justification of Democracy' *Philosophy and Public Affairs* 42: 287-336

-------------------- (2016) 'Political Rule and its Discontents' in David Sobel, Peter Vallentyne and Steven Wall (eds.) *Oxford Studies in Political Philosophy, Volume 2* Oxford: Oxford University Press

Kratzer, Angelika (2012) *Modals and Conditionals* Oxford: Oxford University Press

Kropotkin, Peter (1970) *Kropotkin's Revolutionary Pamphlets: A Collection of Writings by Peter Kropotkin* (ed. Roger N. Baldwin) New York City: Dover

Lackey, Jennifer (2007) 'Why We Don't Deserve Credit for Everything We Know' *Synthese* 158: 345-61

Ladenson, Robert (1980) 'In Defense of a Hobbesian Conception of Law' *Philosophy and Public Affairs* 9: 134-59

Lawford-Smith, Holly (2013) 'Understanding Political Feasibility' *The Journal of Political Philosophy* 21: 243-59

Lehrer, Keith (1976) '"Can" in Theory and Practice: A Possible Worlds Analysis' in Myles Brand and Douglas Walton (eds.) *Action Theory: Proceedings of the Winnipeg Conference on Human Action, Held at Winnipeg, Manitoba, Canada, 9-11 May 1975* Dordrecht: D. Reidel Publishing Co.

Lewis, David (1973) 'Counterfactuals and Comparative Possibility' *Journal of Philosophical Logic* 2: 418-446

Lipsey, Richard G. and Lancaster, Kelvin J. (1956) 'The General Theory of Second Best' *Review of Economic Studies* 24: 11-33

Locke, John (1980) *Second Treatise of Government* ed. C. B. Macpherson. Indianapolis, IN: Hackett Publishing Company

Madison, James (1990) 'The Federalist No. 51' in George W. Carey and James McClellan (eds.) *The Federalist* (266-9) Dubuque, IA: Kendall/Hunt Publishing Company

Majone, Giandomenico (1975) 'On the notion of political feasibility' *European Journal of Political Research* 3: 259-74

Mason, Andrew (2004) 'Just Constraints' *British Journal of Political Science* 34: 251-68

-------------------- (2016) 'Justice, Feasibility, and Ideal Theory: A Pluralist Approach' *Social Philosophy and Policy* 33: 32-54

Mill, John Stuart (2015) *On Liberty, Utilitarianism and Other Essays* (ed. Mark Philp and Frederick Rosen) Oxford: Oxford University Press

Miller, David (1984) *Anarchism* London: J.M. Dent & Sons Ltd.

-------------------- (2008) 'Political philosophy for Earthlings' in David Leopold and Marc Stears (eds.) *Political Theory: Methods and Approaches* Oxford: OUP

Nardin, Terry (2004) 'Justice and Coercion' in Alex J. Bellamy (ed.) *International Society and its Critics* Oxford: Oxford University Press

Nozick, Robert (1974) *Anarchy, State and Utopia* New York: Basic Books

Olsaretti, Serena (2004) *Liberty, Desert and the Market* Cambridge: Cambridge University Press

Pallikkathayil, Japa (2017) 'Persons and Bodies' in Sari Kisilevsky and Martin Stone (eds.) *Freedom and Force: Essays on Kant's Legal Philosophy* Oxford: Hart Publishing

Peter, Fabienne (2012) 'Authority and Legitimacy' in Gerald Gaus and Fred D'Agostino (eds.) *The Routledge Companion to Social and Political Philosophy* London: Routledge

Pettit, Philip (1996) 'Freedom as Antipower' *Ethics* 106: 576-604

-------------------- (1997) *Republicanism: A Theory of Freedom and Government* Oxford: Clarendon Press

-------------------- (2012) *On the People's Terms: A Republican Theory and Model of Democracy* Cambridge: Cambridge University Press

Phillips, Michael (1985) 'Reflections on the Transition from Ideal to Non-Ideal Theory' *Nous* 19: 551-70

Pritchard, Duncan (2012) 'Anti-Luck Virtue Epistemology' *The Journal of Philosophy* 109: 247-79

Räikkä, Juha (1998) 'The Feasibility Condition in Political Theory' *The Journal of Political Philosophy* 6: 27-40

Rawls, John (1993) *Political Liberalism* New York City, NY: Columbia University Press

-------------------- (1999a) *A Theory of Justice: Revised Edition* Cambridge, MA: Harvard University Press

-------------------- (1999b) *The Law of Peoples* Cambridge, MA: Harvard University Press

Raz, Joseph (1979) *The Authority of Law: Essays on Law and Morality* Oxford: Oxford University Press

-------------------- (1988) *The Morality of Freedom* Oxford: Oxford University Press

Renzo, Massimo (2011) 'State Legitimacy and Self-Defence' *Law and Philosophy* 30: 575-601

Ripstein, Arthur (2009) *Force and Freedom: Kant's Legal and Political Philosophy* Cambridge, MA: Harvard University Press

Robeyns, Ingrid (2008) 'Ideal Theory in Theory and Practice' *Social Theory and Practice* 34: 341-62

Ross, Jacob (2012) 'Actualism, Possibilism, and Beyond' in Mark Timmons (ed.) *Oxford Studies in Normative Ethics: Volume 2* Oxford: Oxford University Press

Rossi, Enzo (2012) 'Justice, legitimacy and (normative) authority for political realists' *Critical Review of International Social and Political Philosophy* 15: 149-64

Rossi, Enzo and Sleat, Matt (2014) 'Realism in Normative Political Theory' *Philosophy Compass* 9/10: 689-701

Sangiovanni, Andrea (2012) 'Can the Innate Right to Freedom Alone Ground a System of Public and Private Rights?' *European Journal of Philosophy* 20: 460-9

Sartorius, Rolf (1999) 'Political Authority and Political Obligation' in William A. Edmundson (ed.) *The Duty to Obey the Law: Selected Philosophical Readings* Lanham, MD: Rowman and Littlefield Publishers, Inc.

Schmidtz, David (1990) 'Justifying the State' *Ethics* 101: 89-102

-------------------- (2011) 'Nonideal Theory: What It Is and What It Needs to Be' *Ethics* 121: 772-96

Sen, Amartya (2006) 'What Do We Want from a Theory of Justice?' *The Journal of Philosophy* 103: 215-38

-------------------- (2009) *The Idea of Justice* Cambridge, MA: Bellknap Press

Simmons, A. John (1979) *Moral Principles and Political Obligations* Princeton, NJ: Princeton University Press

-------------------- (1991) 'Locke and the Right to Punish' *Philosophy and Public Affairs* 20: 311-49

-------------------- (1996) 'Philosophical Anarchism' in J. T. Sanders and J. Narveson (eds.) *For and Against the State* Lanham, MD: Rowman and Littlefield Publishers

-------------------- (2001) *Justification and Legitimacy: Essays on Rights and Obligations* Cambridge, UK: Cambridge University Press

-------------------- (2010) 'Ideal and Nonideal Theory' *Philosophy and Public Affairs* 38: 5-36

Simpson, Thomas W. (2017) 'The Impossibility of Republican Freedom' *Philosophy and Public Affairs* 45: 27-53

Singer, Peter (1973) *Democracy and Disobedience* Oxford: Oxford University Press

Smith, M.B.E.(1973) 'Is There a Prima Facie Obligation to Obey the Law?' *The Yale Law Journal* 82: 950-76

Song, Edward (2011) 'Strong and Weak Legitimacy' *Southwest Philosophy Review* 27: 167-74

Sosa, Ernest (2007) *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume I* Oxford: Oxford University Press

Southwood, Nicholas (2016) 'Does "Ought" Imply "Feasible"?' *Philosophy and Public Affairs* 44: 7-45

Stemplowska, Zofia (2008) 'What's Ideal About Ideal Theory?' *Social Theory and Practice* 34: 319-40

-------------------- (2016) 'Feasibility: Individual and Collective' *Social Philosophy and Policy* 33: 273-91

Stemplowska, Zofia and Swift, Adam (2018) 'Dethroning Democratic Legitimacy' in David Sobel, Peter Vallentyne and Steven Wall (eds.) *Oxford Studies in Political Philosophy, Volume 4* Oxford: Oxford University Press

Swift, Adam (2008) 'The value of philosophy in nonideal circumstances' *Social Theory and Practice* 34: 319-40

Valentini, Laura (2009) 'On the Apparent Paradox of Ideal Theory' *Journal of Political Philosophy* 17: 332-55

-------------------- (2011) 'A Paradigm Shift in Theorizing About Justice? A Critique of Sen' *Economics and Philosophy* 27: 297-315

-------------------- (2012a) 'Assessing the global order: justice, legitimacy, or political justice?' *Critical Review of International Social and Political Philosophy* 15: 593-612

-------------------- (2012b) 'Ideal vs. Non-Ideal Theory: A Conceptual Map' *Philosophy Compass* 7/9: 654-664

-------------------- (2012c) 'Kant, Ripstein and the Circle of Freedom: A Critical Note' *European Journal of Philosophy* 20: 450-9

Vallentyne, Peter and Steiner, Hillel (2000) *Left-Libertarianism and its Critics: The Contemporary Debate* Basingstoke: Palgrave

Weber, Max (1948) 'Politics as a Vocation' in H.H. Gerth and C. Wright Mills (eds.) *From Max Weber: Essays in Sociology* London: Routledge and Kegan Paul Ltd.

Wedgwood, Ralph (2009) 'Against Actualism' [Blog post, 11th September 2009] Retrieved from http://peasoup.typepad.com/peasoup/2009/09/against-actualism.html

Wellman, Christopher H. (1996) 'Liberalism, Samaritanism, and Political Legitimacy' *Philosophy and Public Affairs* 25: 211-37

-------------------- (2001) 'Toward a Liberal Theory of Political Obligation' *Ethics* 111: 735-59

-------------------- (2009) 'Rights and State Punishment' *The Journal of Philosophy* 106: 419-39

Wiens, David (2015a) 'Political Ideals and the Feasibility Frontier' *Economics and Philosophy* 31: 447-77

---------------- (2015b) 'Against Ideal Guidance' *Journal of Politics* 77: 433-446

---------------- (2016) 'Assessing ideal theories: Lessons from the theory of second best' *Politics, Philosophy and Economics* 15: 132-49

---------------- (2017) 'Will the Real Principles of Justice Please Stand Up?' in Michael Weber and Kevin Vallier (eds.) *Political Utopias: Contemporary Debates* Oxford: Oxford University Press

Williams, Bernard (2005) 'Realism and Moralism in Political Theory' in his *In the Beginning was the Deed: Realism and Moralism in Political Argument* Princeton, NJ: Princeton University Press

Woodcock, George (1963) *Anarchism* Harmondsworth, Middlesex: Penguin Books

Wolff, Jonathan (1991) *Robert Nozick: Property, Justice and the Minimal State* Stanford, CA: Stanford University Press

Wolff, Robert Paul (1970) *In Defense of Anarchism* New York City: Harper Torchbooks

Wright, Erik Olin (2006) 'Compass Points: Towards A Socialist Alternative' *New Left Review* 41: 93-124

Wyckoff, Jason (2010) 'The Inseparability Thesis: Why Political Legitimacy Entails Political Obligations' *Southwest Philosophy Review* 26: 51-9

Zhu, Jiafeng (2017) 'Farewell to Political Obligation: In Defense Of A Permissive Conception of Legitimacy' *Pacific Philosophical Quarterly* 98: 449-69