

The Moral Tyranny Constraint

Tyrants: for centuries and centuries, you have sucked our blood. The tears which you have made us spill would be enough to drown you From today forward, there will not be a man who dares to make others obey him; there will not be a man who exploits the work of another man Comrades: we must complete social justice. Let us cut off the head of the hydra and take possession of all that exists for the well-being of all. Long live Land and Liberty!

Ricardo Flores Magón, *Land and Liberty*

A recurrent theme in libertarian thought is that persons should not be allowed to discretionarily impose costs upon others. For example, in an influential polemic, William Sumner objects to publicly funded policing of vagrancy on the grounds that “the industrious workman going home from a hard day’s work . . . is mulcted of a percentage of his day’s earnings to hire a policeman to save the drunkard from himself” (1918, 480). This objection is echoed by Ludwig von Mises’ complaint that social insurance for farmers entails that “if they blunder . . . the government forces the consumers, the taxpayers, and the mortgagees to foot the bill” (1998, 583). Similarly, Murray Rothbard objects to state-imposed egalitarian redistribution because “others are being forced to pay the cost” of helping the poor (1995, 53). More generally, Eric Mack argues that people must not be granted the “moral liberty to subordinate us to their purposes, that is, to impose sacrifices upon us to advance their ends” (2010, 60).¹ And Jason Brennan has recently argued that democracy is problematic because:

[In a democracy] some people impose their decisions on others. If most voters act foolishly, they don’t just hurt themselves. They hurt better-informed and more rational voters, minority voters . . . and foreigners who are unable to vote

¹ Mack takes this requirement to follow from the twin libertarian commitments of (a) taking seriously the separateness of persons and (b) affirming that one is “allowed to live one’s life in one’s own chosen way” (2010, 60).

but still are subject to or harmed by that democracy's decisions If the majority makes a capricious decision, others have to suffer the risks." (2016, 9)

The shared presumption of these claims is that there is something objectionable about a social system that allows some people to impose costs on others, either through negligence or malicious intent. The purpose of this book is to explore the consequences of taking this presumption seriously. It will contend that those who are genuinely committed to this libertarian presumption should ultimately endorse the egalitarian variety of anarchism presented in Chapter 1. This chapter will lay the groundwork for this argument by (1) providing a formal statement of the libertarian presumption, (2) explicating that statement, (3) defending its plausibility, and (4) explaining how it both entails two of the core libertarian principles introduced in Chapter 1 and supports a luck egalitarian approach to distributive justice.

2.1 The Moral Tyranny Constraint

To begin, note that the libertarian presumption is best understood as a constraint upon which moral theories count as acceptable. If it is wrong to impose costs upon others, then a moral theory that licenses such cost imposition will be extensionally inadequate. Thus, a moral theory is acceptable only if it does not license such imposition. For these purposes, this general conditional proposition will be restated a bit more narrowly so as to make it as uncontroversial as possible. With the addition of various qualifiers, the posited constraint will be able to sidestep potential counterexamples while still delivering the promised libertarian principles and egalitarian anarchist conclusion introduced in Chapter 1. Call this restatement of the libertarian presumption the *moral tyranny constraint*:

The Moral Tyranny Constraint – A theory of duties is acceptable only if full compliance with that theory would not allow any person to unilaterally, discretionarily, and foreseeably act in a way that would leave others with less advantage than they would have possessed had the agent made some other choice.

There are quite a few qualifying terms packed into this constraint, each of which will be explicated in the subsequent section. Before discussing these details, however, it is worth clarifying the constraint by describing the two ways in which it might be violated by a moral theory. First, a theory violates the constraint if an agent is able to carry out some action that (unilaterally, discretionarily, and foreseeably) worsens another's

position and the theory does not give the latter a claim to full compensation. For example, if a theory does not give persons a claim to redress for harm-inflicting actions like assault, then it would run afoul of the constraint. In other words, the constraint is violated by a theory that permits agents to carry out a set of actions that would collectively leave others worse off. Second, a theory violates the constraint if it assigns to any person *P* a Hohfeldian power to oblige another person *Q* to do some action ϕ where *Q* ϕ -ing would leave *Q* (or some third-party) worse off – for example, if it holds that *P* can oblige *Q* to destroy *Q*'s favorite painting. While the mere imposition of the obligation would not leave *Q* worse off in this case, her compliance with the imposed duties would; thus, the theory violates the moral tyranny constraint.

2.2 Explicating the Constraint

This section will explicate the moral tyranny constraint's many qualifying terms, thereby helping to clarify the constraint and give it determinate content. Specifically, there are six components of the constraint that are in need of explication. First, note that the constraint only applies to theories of duties, that is, theories that assign to every action some deontic status such that the action is declared to be either permissible, impermissible, or obligatory.² More precisely, the constraint and the subsequent arguments of the book presume something proximate to Hohfeld's (1913) schema of deontic incidents. According to this schema, *duties* (or, alternatively, *obligations*) are taken to entail the existence of correlative *rights* (or, alternatively, *claims*) possessed by others.³ A *permission* to do a thing (or, alternatively, a *privilege*) is possessed by a person when no one has a claim that she not do that thing; that is, all others have a correlative *no-claim* with respect to her doing that action.⁴ A person possesses a *power* when she is able to alter her own or others' incidents, where the correlative of a

² Permissible actions might be further divided into those that are permissible *tout court*, those that are supererogatory (Urmson 1958), and those that are suberogatory (Driver 1992). Such theories of duties contrast with *aretaic theories*, which take the primary moral judgment to be assessments of character rather than assessments of actions.

³ As noted in Footnote 43 of Chapter 1, there may also be *non-directed duties* that do not entail a correlative claim. In other words, people who possess these duties would still be obliged to carry out certain actions, but they would not *owe* such actions to any particular person(s). For a recent critical discussion of the directedness of duties, see Rowan Cruft (2019).

⁴ Some people also call permissions "liberties"; however, others such as Thomson use the term "liberty" to refer to the conjunction of a permission to do some action and a claim against all others that they not interfere with that action (1990, 53–4). Thus, the following argument will avoid using this language and will, instead, generally use the term "permission."

power is a *liability*. For example, if one person is able to impose a duty on another, the former has a Hohfeldian power while the latter has a liability. Finally, when someone lacks a power to alter an incident, they have a *disability* vis-à-vis that incident, with the possessor of the protected incident having an *immunity* from having that incident negated.

Second, the moral tyranny constraint is stated in terms of diminishing others' *advantage*. "Advantage" here should be taken to have the same referent as the term "advantage" that appeared in Section 1.6; that is, it should be understood to not have any specific content but, rather, function as a placeholder for whatever one takes to be the relevant currency of distributive justice. For example, one might think that what matters is whether a person's welfare is diminished. Alternatively, one might think that the relevant question is whether she is left with fewer goods, where there is some objective list of goods. Or, perhaps, one must consider the output of some function that takes as its arguments a person's objective goods, her welfare, and/or some other property she possesses.⁵ Because the argument of the book is compatible with any of these proposals, it will remain neutral regarding which one is best, with the term "advantage" referring to any favored currency of well-being. However, for the book's argument to be valid, the term must have a consistent referent throughout (i.e., it must refer to the same thing when it appears in the moral tyranny constraint as it does when it appears in the anarchist conclusion presented in Section 1.6).

Third, when the constraint asserts that agents must not be able to leave others worse off under conditions of full compliance – or, to introduce a bit of simplifying terminology leave others worse off_{FC}/with less_{FC} – it is making a counterfactual claim rather than a temporal one. In other words, when asking whether *P* is able to act in a way that would leave *Q* worse off in the full-compliance world, the question is not whether there is some action ϕ that *P* can take such that the combination of her ϕ -ing and full compliance causes *Q* to have less advantage than she had before *P* ϕ -ed. Rather, the question is whether there is some alternative action ψ where (i) ψ -ing is impossible with *P* ϕ -ing (e.g., ψ -ing might simply be identical to the omission of not ϕ -ing) and (ii) *Q* would have more advantage if *P* were to ψ under conditions of full compliance than she would if *P* were to ϕ under such conditions.

⁵ Cohen uses the term "advantage" to refer to some combination of both welfare and resources (2011, 18). However, the term used here should not be understood to be co-referential with Cohen's notion.

The fourth point of clarification pertains to the constraint's qualification that a theory of duties must not license a person to *unilaterally* leave others worse off/with less advantage under conditions of full compliance. The "unilaterally" qualifier should be understood as follows: *P* unilaterally leaves *Q* worse off_{FC} by ϕ -ing if and only if (a) *Q* is unable to avoid being left worse off_{FC} once *P* has ϕ -ed and (b) *Q* does not consent to being left worse off in this way. To illustrate, suppose that a moral theory assigns *P* the power to impose a conditional obligation on *Q* such that *P* can make it obligatory that if *Q* ϕ -s then *Q* ψ -s. Further, suppose that *Q* would have just as much advantage if she ϕ -ed as she would if she did not ϕ , but would be left worse off if she were to ψ relative to her not ψ -ing. In this case, Condition (a) is not satisfied, as *Q* could discharge her obligation by simply not ϕ -ing (as making the antecedent of the conditional false renders the entire conditional true) and be no worse off as a result.⁶ Similarly, if *P* can oblige *Q* to ψ – but *P* acquires that power only if she first receives *Q*'s consent – then Condition (b) would not be satisfied and, thus, *P* would not be able to unilaterally leave *Q* with less_{FC}.

This qualification is included because there is seemingly nothing defective about a theory that permits the imposition of costs that are voluntarily accepted. Indeed, almost all rights theorists would affirm that it is permissible to impose costs upon a person if she consents to that imposition. Granted, the proposed qualification takes an expansive view of what qualifies as voluntary, as it treats a person's response to some action as implying the voluntary acceptance of the consequences of that response. Some might think that this is too permissive and insist that the analysis of "unilaterally" ought to be broadened such that the set of unilateral costs excludes only those costs that were consented to – that is, Condition (a) ought to be removed from the proposed analysis leaving only Condition (b). However, such an adjustment would make the moral tyranny constraint more stringent than it needs to be. As will be made clear in Section 2.5, the moral tyranny constraint still entails two consequential and demanding libertarian principles even assuming the original, narrower specification of what counts as "unilateral." Thus, the proposed constraint is qualified in a way that errs on the side of modesty, avoiding controversy

⁶ Note that if not ϕ -ing were more costly to *Q* than ϕ -ing, then *P* *would* be able to unilaterally leave *Q* worse off_{FC} by imposing the conditional obligation in question. Thus, while the moral tyranny constraint allows that *P* might impose conditional obligations on *Q*, it still imposes strict limits on which conditional obligations *P* can impose. Specifically, there must always be some option available to *Q* such that she ends up no worse off_{FC} than if the obligation had not been imposed.

by counting more theories as acceptable rather than fewer (while still generating philosophically significant results).

Fifth, the proposition that a person *discretionarily* leaves others worse off_{FC} should be understood as asserting two things. First, it asserts that there was an alternative choice that the agent could have made that would not have resulted in the person ending up worse off_{FC}. Second, this alternative choice has to be not merely modally available in the just-mentioned sense but also *morally* available in the sense that the agent has no duty of justice to refrain from making that choice. For example, suppose that a doctor could substantially improve a patient's life by performing an invasive medical procedure to which the patient has refused consent. If the doctor chooses not to perform the procedure, she will leave the patient with worse off_{FC} than if the procedure were performed. However, there is seemingly nothing problematic about a moral theory that countenances this result – that is, that denies that the patient has any claim to compensation. This is because the doctor is simply doing what she *has* to do when she chooses not to carry out the procedure (in the normative sense of “has”). Thus, a theory only seems intuitively unacceptable when it licenses people to act in non-obligatory, cost-imposing_{FC} ways.⁷

Finally, a moral theory only seems unacceptable if it countenances people *foreseeably* leaving others worse off_{FC}. For example, consider the case of a hiker who falls into an abandoned well due to the mouth being hidden by moss and leaves. A moral theory that assigns a bystander an obligation to rescue the hiker – where the rescue would be moderately costly to the bystander – does not seem obviously defective. By contrast, if the theory obliged the bystander to rescue a spelunker who decided to

⁷ This case draws attention to the fact that the moral tyranny constraint *must* include the “discretionarily” qualifier lest it entail that there are conflicting duties in any case where an obligatory action does not maximize others' advantage. To see why, consider the general case where Q has a claim that $P \phi$, where $P \phi$ -ing leaves Q with less advantage than if P did not ϕ . What happens if $P \phi$ -s? If the moral theory in question does not assign Q some sort of claim to compensation from P , then it would violate the moral tyranny constraint, as P 's ϕ -ing would leave Q worse off_{FC} (as full compliance will not offset Q 's lost advantage). Thus, Q must have a claim to compensation if $P \phi$ -s. But what would ground such a claim given that P is merely discharging her duty to Q by ϕ -ing? Seemingly, Q has a claim to compensation only if P 's ϕ -ing infringed on some claim of Q 's that P not ϕ ; that is, P has both a duty to ϕ and a duty to refrain from ϕ -ing. While such a conflict of rights is not a logical contradiction, it is an undesirable thing for a theory of duties to affirm. It would therefore be a significant problem if the moral tyranny constraint entailed that there was a conflict of duties in any case where someone was obliged to act in a way that did not maximize others' advantage. Thus, the “discretionarily” qualifier must be included in the constraint to avoid this implication.

explore the well despite knowing the risks involved, that would seem to be a clear instance of extensional inadequacy. Given that this apparent difference is best explained by the unforeseeability of the former outcome, the moral tyranny constraint must be qualified such that it only rules out theories that license people foreseeably leaving others worse off_{FC}.

When a moral theory meets all of the aforementioned conditions – that is, when it affirms that a person can unilaterally and discretionarily act in a way that leaves others foreseeably worse off_{FC} – it seemingly institutes an unacceptable sort of moral tyranny. The hallmark of a tyrant is that she is able to impose discretionary costs on her subjects without any sort of legal restraint. Analogously, a theory that violates the moral tyranny constraint allows people to willfully diminish others' well-being without any sort of normative restraint. It is the licensing of such behavior that the libertarians cited in this chapter's opening paragraph seemingly find so objectionable. Each of their quoted complaints represents an objection to legal systems that license the imposition of costs on others, where such objections seemingly entail a correlative objection to any theory of duties that licenses cost-imposing actions. Thus, the moral tyranny constraint can be understood as giving more precise expression to this persistent libertarian complaint.

2.3 Defending the Constraint

The introduction of this chapter noted that numerous libertarians presume something proximate to the moral tyranny constraint. This section will provide an argument in defense of this libertarian presumption. Specifically, it will present three reasons for accepting the moral tyranny constraint that will appeal to both libertarians and non-libertarians alike. It thereby aims to show that the moral tyranny constraint is an attractive and independently plausible meta-principle.

First, the constraint can be seen as following from the *separateness of persons argument* that liberals advance against utilitarianism.⁸ Utilitarians hold that a cost can be permissibly imposed upon one person if the imposition of that cost is a necessary and sufficient condition for providing greater benefit to another. In support of this claim, utilitarians will often present something like the following *argument from prudential choice*. They begin by noting that humans regularly impose costs upon themselves for

⁸ Notably, versions of this argument have been presented by both John Rawls (1971, 26–7) and Nozick (1974, 32–4), though the two draw very different conclusions about what follows from the separateness of persons (beyond the negation of utilitarianism).

the sake of obtaining greater benefits in the future – a practice that seems not only morally unproblematic but rationally demanded. Further, they maintain that if there is nothing problematic about a person imposing costs upon herself for greater future benefit, then there is nothing problematic about imposing costs on one person to provide greater benefit to another. Thus, utilitarians conclude that there is nothing problematic about sacrificing one person's well-being for the sake of providing greater benefit to another.⁹

The aforementioned separateness of persons argument is best understood as an objection to this argument from prudential choice. Specifically, it is an objection to the second premise of the utilitarian's argument – that is, the premise that, if it is permissible for a person to impose costs on herself for greater future benefit, then it is permissible to impose costs on one person in order to provide greater benefit to another. The fact that persons are separate rather than some unified social creature renders this inference implausible. The separateness of persons objection points out that there is a crucial disanalogy between prudential sacrifice and utilitarian sacrifice that prevents the permissibility of the former from implying the permissibility of the latter.¹⁰

⁹ J. J. C. Smart makes this argument explicitly in reply to Rawls' (1958) suggestion that there is something unfair about imposing costs on some to maximize overall utility:

[I]f it is rational for me to choose the pain of a visit to the dentist in order to prevent the pain of toothache, why is it not rational of me to choose pain for Jones, similar to that of my visit to the dentist, if that is the only way in which I can prevent a pain, equal to that of my toothache, for Robinson?

(Smart and Williams, 1973, 37)

¹⁰ This way of explicating the separateness of persons argument heads off a potential reply suggested by Mack (2018) in his elaboration of Nozick's version of the argument. Mack posits that the utilitarian might argue that the reason that a person can permissibly impose costs on herself for future benefits is that one ought to maximize utility. As he puts it, "we start with the unrestricted rationality of minimizing costs (or maximizing benefits); and the principle of individual choice is simply the application of that principle of social choice to the special case in which there is only one agent" (2018, §2.2). However, given the chapter's proposed interpretation of the separateness of persons argument, this reply would beg the question, as the general principle that one ought to maximize benefits irrespective of whether there is one person or many is the proposition that has to be demonstrated. Thus, it cannot be assumed as a starting premise.

An additional advantage of stating the argument in this way is that it allows the appeal to separateness of persons to qualify as a supporting argument for liberal deontological positions (as opposed to a mere restatement of those positions). A common tendency in the recent literature is to treat the proposition that persons are separate as merely an alternative way of expressing some other deontological commitment such as the thesis that persons are owed respect and, thus, cannot be treated as a mere means (Zwolinski 2008, 150–2) or the contention that persons have moral authority over their own lives that trumps moral reasons to promote the common good (Mazor 2019, 192–3). However, such an approach strips the separateness of persons argument of its dialectical force. In the debate between utilitarians and deontologists, the former assert that you ought to maximize utility even if that requires treating someone as a mere means or denying her

Proponents of the separateness of persons argument conclude that one cannot permissibly impose costs on some people to benefit others. Admittedly, this conclusion does not strictly follow from the separateness of persons argument, as the utilitarian's conclusion might still be true even if the separateness of persons objection renders her argument from prudential choice unsound. However, one might think that the argument from prudential choice is the only plausible way of justifying something that is *prima facie* unjustified. Thus, if that argument fails due to the separateness of persons objection, then it follows that it is not permissible to sacrifice some to benefit others.

If one accepts this conclusion, then one should also endorse the moral tyranny constraint. Notably, the objection's conclusion condemns utilitarian sacrifice even given the fact that the provision of benefit is at least a plausible candidate for justifying cost imposition. By contrast, there will be many cases of discretionary cost imposition where this is no countervailing moral consideration that might justify the imposed cost. Given that utility-maximizing cost imposition is not permissible even given countervailing moral considerations, it seemingly follows from the separateness of persons argument that discretionary cost imposition (which lacks this justificatory advantage) is also impermissible.

Compare this result with the moral tyranny constraint, which holds that it is not permissible to unilaterally, foreseeably, and discretionarily impose costs_{FC} on a person. The crucial difference between this claim and the conclusion of the separateness of persons argument is that the latter condemns the imposition of actual-world discretionary costs while the former condemns the imposition of full-compliance-world discretionary costs. However, those who hold that it is impermissible to impose actual-world costs are seemingly committed to affirming that it is impermissible to impose full-compliance-world costs. To see why, consider the following case: if P ϕ -s, she leaves Q worse off; by contrast, if P ψ -s, she does not leave Q worse off but *does* leave her worse off_{FC}. Further, assume that both actions are discretionary. According to the conclusion of the separateness of persons argument, P 's ϕ -ing would be impermissible because it imposes

moral authority over how her life goes; by contrast, deontologists assert that you ought not maximize utility under such circumstances. If affirming the separateness of persons merely expresses the proposition that one ought not treat persons as a means and/or individuals' moral authority trumps promoting the common good, then it merely reasserts the deontological position rather than providing a reason for favoring that position over the utilitarian one. By contrast, the proposed separateness of persons argument does provide such a reason.

a cost on Q – where ϕ -ing imposes a cost on Q if and only if Q 's loss of advantage is a direct function of P ϕ -ing. In other words, had P not ϕ -ed, then Q would not have lost advantage and *for this reason* P 's ϕ -ing is impermissible.

Now, consider the world where P ψ -s and everyone complies with their moral requirements. In such a world, P 's action stands in an identical counterfactual relation to Q 's loss of advantage. Thus, one might conclude that P 's ψ -ing is also impermissible. However, this would be a mistake, as it is actually P 's ψ -ing *conjoined* with certain acts of compliance that lead to Q being worse off; absent such compliance, P ψ -ing would not leave Q worse off relative to the world where P does not ψ . In other words, the appropriate conclusion to draw is that *the conjunction of P ψ -ing and the acts of compliance* is impermissible. Given that P ϕ -ing is impermissible because Q 's worsened position is a direct function of P ϕ -ing, the fact that Q 's worsened position is also a direct function of P ψ -ing plus others' acts of compliance entails the impermissibility of that set of actions.¹¹ This, in turn, implies that if these other acts are all permissible, then P ψ -ing is impermissible. Further, note that the world of full compliance is, by definition, a world where all other persons act *permissibly*. Thus, it must be the case that P ψ -ing is impermissible. In other words, if one accepts the conclusion of the separateness of persons argument – that is, that it is impermissible to discretionarily leave others worse off – then one should also accept the moral tyranny constraint's implication that it is impermissible to discretionarily leave others worse off_{FC}.

A second reason for endorsing the moral tyranny constraint is that it is a less demanding – and, thus, less controversial – version of the popular neo-republican rejection of domination advanced by Philip Pettit (2012). Pettit contends that there is something morally objectionable about a state of affairs where one person has the ability to limit another's freedom by removing one of her options. He suggests that the presence of such

¹¹ One might worry that this argument departs from some of the previous discussion by putting deontic propositions in impersonal terms – that is, there is some action-including state of affairs that is obligatory/permissible – rather than in agential terms such that some agent is said to be obliged/ permitted to do a thing. However, first, one could seemingly frame the entire argument of the book in impersonal terms without issue. Any inconsistency in deontic language merely reflects localized stylistic choices that aid in exposition. Alternatively, one might take there to be some way of bridging the agential-impersonal gap, for example, by holding that a person is obliged to act in some way if and only if it is obligatory that the person act in that way. This presumption is proximate to what has become known as the “Meinong-Chisholm Reduction” after two early proponents of this view (namely, Alexius Meinong (1972) and Roderick M. Chisholm (1964)). However, this view is not without critics (see, e.g., Jacob Ross (2010)).

domination undermines egalitarian relations between persons, precluding their ability to “look others in the eye without . . . fear or deference that a power of interference might inspire; [to] walk tall and assume the public status . . . of being equal” (2012, 84). Additionally, beyond this relational egalitarian concern, the objection to domination might be viewed as a natural extension of a more primitive concern with freedom, where the restriction of freedom – and, more strongly, people having the ability to restrict others’ freedom – is taken to be objectionable.

To see why the moral tyranny constraint is a weaker version of the republican thesis, it will be helpful to restate the latter in deontic terms. Specifically, a Pettit-influenced republican holds that it is impermissible to realize or preserve a state of affairs where one person has the ability to remove options from another’s option set. By contrast, the moral tyranny constraint limits which options it is permissible to remove. Specifically, it entails that an agent is forbidden from carrying out a conjunction of actions that would remove advantage-preserving_{FC} options such that someone is left worse off_{FC}. For example, if (a) P ϕ -ing and then ψ -ing would leave Q worse off_{FC} and (b) P ϕ -s, then the moral tyranny constraint would imply that it is impermissible for P to ψ ; that is, the moral tyranny constraint implies that it is impermissible for P to remove an option from Q ’s option set by ψ -ing (namely, the option where Q does her most-preferred action and Q does not ψ). This makes the moral tyranny constraint weaker than the standard republican position in two respects. First, it merely forbids *option removal* rather than forbidding people from having the *capacity* to remove options (where someone lacking the capacity to remove an option entails that they will not remove that option but where the converse of this conditional is false). Second, the republican thesis is concerned with the removal of all options while the moral tyranny constraint forbids only the removal of advantage-preserving_{FC} options, which are a proper subset of all options.¹² Thus, the satisfaction of the republican principle will entail the satisfaction of the moral tyranny constraint but not vice versa. Given this relation, republicans critical of domination should also accept the moral tyranny constraint both because it seems to capture some of the motivating concerns expressed in the previous paragraph and because there is no reason to reject a principle

¹² It is assumed here that changing the moral status of some option via the imposition of an obligation also counts as option removal. For example, if P obliges Q to stay out of the park after midnight, P removes an option from Q , namely, the option of going in the park after midnight while discharging her duties. In other words, options are assumed to be individuated in a fairly fine-grained fashion such that actions with different deontic statuses represent different options.

that is fully satisfied if one's own principle is satisfied. Additionally, because the moral tyranny constraint is less stringent than the standard republican position, it is vulnerable to fewer objections (e.g., that there are cases where having a mere *capacity* to remove some arbitrary option is morally unproblematic).

Finally, one ought to accept the moral tyranny constraint because it would be the result of a process of reflective equilibrium. Most famously championed by John Rawls (1971), the practice of reflective equilibrium involves rendering one's set of normative beliefs coherent, where this set includes one's particular moral judgments, the general moral principles that support those judgments, and the theoretical *desiderata* (i.e., meta-principles) that determine which principles are acceptable. As noted in Section I.2, coherence is an important theoretical virtue of any given normative position, where genuine coherence requires that relations of logical entailment obtain between that position's various propositions. In other words, if one embraces the method of reflective equilibrium – with the associated presupposition that coherence is a theoretical virtue – then one has reason to accept moral principles that entail a large number of one's accepted particular moral judgments or moral principles. Similarly, one has reason to accept those theoretical *desiderata* that entail a large number of one's accepted general principles.

In further defense of accepting theoretical *desiderata* (or moral theories) because they entail many accepted moral theories (or particular judgments), one might draw an analogy between the virtues of normative theories and the virtues of scientific theories. Philosophers of science typically maintain that explanatory power is a virtue of scientific theories: One has more reason to accept a theory that explains a large number of observed phenomena than one that explains fewer, *ceteris paribus*.¹³ Analogously, one might think that general normative principles (theoretical *desiderata*) stand in an explanatory relation to particular moral judgments (general moral principles).¹⁴ Together, these claims would entail that one has reason to accept a moral principle (theoretical *desideratum*) in proportion to the number of accepted particular judgments (moral principles) it entails.

If this is right, then there is further reason for accepting the moral tyranny constraint, as it entails a number of influential moral principles.

¹³ For a defense of this point, see Bas van Fraassen (1980, 98).

¹⁴ For a sustained argument that normative theories are, in important ways, analogous to scientific explanations, see Jesse Spafford (2021a).

Specifically, the next section will argue that the Lockean proviso, the consent theory of legitimacy, and luck egalitarianism's incorporation of responsibility all follow from the moral tyranny constraint.¹⁵ It will thereby show that the moral tyranny constraint has both significant explanatory power – that is, it explains why these various principles obtain – and helps to establish the coherence of a number of attractive views. This makes the moral tyranny constraint a strong candidate theoretical *desideratum* to include in any reflective equilibrium.

In sum, there are a number of reasons for adopting the moral tyranny constraint. Two are foundationalist in the sense that the reason for accepting the constraint is that it follows from some other plausible position, for example, the separateness of persons argument or the republican critique of domination. The third is coherentist in the sense that the reason for accepting the moral tyranny constraint is that the constraint entails – and thereby helps to explain – various other attractive positions (as will be discussed in the subsequent section). Thus, irrespective of which sort of justification one favors, one has reason to endorse the moral tyranny constraint.

2.4 Three Implications of the Constraint

This section will argue that three of the positions introduced in Chapter 1 – namely, the consent theory of legitimacy, the Lockean proviso, and luck

¹⁵ Additionally, note that the various claims advanced in the opening paragraph of this chapter also follow from the constraint. For example, von Mises (1998) can be understood as denying that social insurance for farmers is just, as that entails that farmers who make imprudent decisions are *entitled* to the transfers they receive and, thus, can leave others worse off_{FC} when they make such decisions. Similar remarks apply to Brennan's (2016) objection to democracy. If democracy were a just system of government, then hooligans who vote for dangerous policies would have a *right* that those policies be implemented. This, in turn, would imply that these hooligans could, by assembling a simple majority, leave others worse off_{FC}. In both cases, the solution is to propose a moral theory that does not license some people to impose costs_{FC} upon others. While this does not stop imprudent farmers or hooligan voters from leaving others worse off as a matter of empirical fact, it does prevent them from leaving others worse off_{FC}, which is all that the moral tyranny constraint requires (and all that should really concern those doing normative theorizing). Admittedly, it is unlikely that non-libertarians will find Brennan's or von Mises' positions attractive. For such skeptics, the fact that the moral tyranny constraint entails a rejection of democracy or social insurance will count *against* including the constraint in the ultimate reflective equilibrium. However, any coherent normative position will likely entail some unfortunate conclusions; thus, the question is whether its desirable implications outweigh the undesirable ones and how this net assessment compares to the assessment of other rival positions. The hope is that non-libertarians will find the implications discussed in the subsequent section to be more attractive (particularly the luck egalitarian implication) to the point where they might be willing to accept the moral tyranny constraint even if that means that they have to accept some of its less attractive implications as a consequence.

egalitarianism's responsibility component – follow from the moral tyranny constraint. To begin, consider the consent theory of legitimacy's assertion that a state or person can impose obligations on others (if and) only if the latter have consented to being morally bound in this way. While there have been many arguments presented in defense of this position (see, e.g., Simmons 2001), it also follows from the moral tyranny constraint's contention that moral theories cannot license persons to unilaterally, discretionarily, and foreseeably leave others worse off_{FC}. Recall from Section 2.1 that one way in which a person might leave others worse off_{FC} is by imposing obligations on them that would be costly to discharge. Thus, the constraint entails that any power to impose such obligations must be restricted such that the imposition cannot be carried out foreseeably, discretionarily, and unilaterally.

This result entails that the moral tyranny constraint is incompatible with any theory of legitimacy that does not have consent as its necessary condition. As noted in Section 1.1, a legitimate authority has the power to oblige others via the issuing of edicts, where this power is *content independent* in the sense that the authority can oblige others to act in some way irrespective of the properties of that act (with the possible exception of the act being morally prohibited). Thus, a legitimate authority has the power to impose obligations irrespective of whether or not the obligor would be worse off if she discharged the obligation; that is, she has the ability to leave others worse off_{FC}. Further, there is seemingly no way to either eliminate the discretionary character of legitimacy or make it such that legitimate authorities can impose only those obligations that do not foreseeably leave others worse off_{FC}. Note that the latter restriction would negate the content-independent character of legitimacy: The proposition that an agent can impose only those obligations that do not foreseeably leave obligors worse off_{FC} is just the kind of content-based restriction that legitimacy lacks as a matter of definition. Similarly, limiting legitimacy such that it becomes a power to *non-discretionarily* impose costly obligations (i.e., impose such obligations only when there is no other permissible option available) strips the power of its essential character. While perhaps there may be such a power, it would only loosely resemble the power that is at issue when philosophers debate whether states are legitimate.

Given that legitimate authorities necessarily have the power to discretionarily and foreseeably impose costly obligations on others, there is only one way to make the power of legitimacy compatible with the moral tyranny constraint: Make others' consent a necessary condition for legitimacy obtaining. If consent is a necessary condition of legitimacy, then

legitimate authorities cannot *unilaterally* impose costly obligations on others, as those others will be able to fully control whether the authority leaves them worse off_{FC}. By contrast, the absence of this necessary condition entails that legitimate authorities can unilaterally, foreseeably, and discretionarily leave others worse off_{FC}.¹⁶ Thus, the moral tyranny constraint entails the consent theory of legitimacy.¹⁷

A similar argument can be given for why the Lockean proviso – or, more precisely, its posited necessary condition of initial appropriation – follows from the moral tyranny constraint.¹⁸ Note that, much like legitimate authorities, those who appropriate natural resources and convert them into private property impose obligations on others (namely, obligations to refrain from using or making nonconsensual contact with the appropriated thing).¹⁹ Thus, the moral tyranny constraint entails that one must posit some necessary condition of initial appropriation to preclude appropriators from unilaterally, discretionarily, and foreseeably leaving others worse off_{FC} via this obligation imposition. One option is to adopt the consent theorist's approach and make consent a necessary condition of appropriation, thereby precluding appropriators from *unilaterally* leaving others worse off_{FC}. However, almost all proponents of initial appropriation reject this option on the grounds that it is too stringent. Given the difficulties of

¹⁶ As noted in Section 2.2, it is possible to nonconsensually and non-unilaterally impose costly obligations on others if those obligations have the right sort of conditional structure. However, an authority whose normative power is limited to imposing such obligations cannot qualify as a *legitimate* authority, as such a restriction contradicts the content independence that is an essential characteristic of legitimacy. Only consent can render content-independent obligation imposition non-unilateral.

¹⁷ This is a slight overstatement of what has been demonstrated. Strictly speaking, there is a stronger and weaker version of consent theory, where the former holds that consent is a necessary *and sufficient* condition of legitimacy while the latter holds that it is merely a necessary condition. The preceding argument shows only that the weak version of consent theory follows from the moral tyranny constraint. One might thereby endorse both the moral tyranny constraint and hold that even consent does not allow authorities to impose obligations on others via edict (perhaps for reasons such as those advanced by Robert Paul Wolff (1970)). However, insofar as one takes consent to be a promising ground for obligation imposition (e.g., in the case of promissory obligations), then one should take consent to be a sufficient condition of legitimacy – and, thus, accept the strong version of consent theory in light of the moral tyranny constraint's implication that consent is also a necessary condition of legitimacy. For a quick argument along these lines against Wolff's position, see Simmons (1987, 269fn2).

¹⁸ Note that there are also stronger and weaker versions of the proviso analogous to the stronger and weaker versions of consent theory discussed in the previous footnote. And, just as was true of consent theory, the moral tyranny constraint implies only the weaker version of the proviso that makes non-worsening a necessary – but not sufficient – condition of successfully appropriating unowned resources.

¹⁹ Much more will be said about the relationship between legitimacy and initial appropriation in Chapter 4.

obtaining universal consent, this standard would unacceptably preclude the establishment of any private property.²⁰ Thus, a less controversial way to satisfy the moral tyranny constraint is to permit unilateral appropriation but hold that such appropriation cannot leave others worse off_{FC} – that is, affirm Chapter 1’s interpretation of the Lockean proviso.

As was noted in Section 2.2 of that chapter, the Lockean proviso is typically presented as asserting that an act of initial appropriation succeeds if and only if it does not leave anyone worse off *tout court*. However, it was argued there that this interpretation of the proviso is unacceptable because it is trivially satisfied by every act of appropriation. Thus, the section concluded that the proviso is better understood as asserting that it is *full compliance with the established claims* that must not leave others worse off (where this thesis can now be stated using the “worse off_{FC}” shorthand presented earlier). What is now hopefully apparent is that this adjustment makes the proviso an application of moral tyranny constraint to the appropriation of private property: The constraint insists that moral theories cannot allow persons to leave each other worse off_{FC}, and the proviso makes it such that the posited theory of property rights complies with this restriction.

Further, one can now see that the moral tyranny constraint also entails the final adjustment that Section 1.2 made to the Lockean proviso. The proposal there was that it is not just appropriators’ established claims that must not leave others worse off_{FC}; rather, any potential waiving of the established claims must also not leave others worse off_{FC}. The justification for this adjustment appealed to a premise about arbitrariness: Given that the power to establish property claims is constrained by a non-worsening condition, it seems unacceptably arbitrary to not impose this same constraint on the power to waive these claims. However, now that the moral tyranny constraint has been introduced, the appeal to arbitrariness is no longer needed, as the constraint directly entails that people cannot have any normative power that enables them to leave others worse off_{FC}. Given that initial appropriation entails the existence of powers to waive the established claims, it follows that initial appropriation can succeed only if any subsequent exercise of these powers would not leave anyone worse

²⁰ See van der Vossen (2019, § 3) and Mack (2010). For an early and influential rejection of consent as a necessary condition of appropriation, see Locke (2005, §28). That said, Chapter 4 will raise some complications for those who want to reject consent as a necessary condition of appropriation without accepting the broader moral tyranny framework presented in this book.

off_{FC}. In this way, the moral tyranny constraint entails the revised Lockean proviso presented in Section 1.2.²¹

So far it has been argued that two prominent libertarian theses follow from the moral tyranny constraint: the consent theory of legitimacy and the (slightly adjusted) Lockean proviso. However, there are also influential non-libertarian positions that follow from the moral tyranny constraint. Most notably, luck egalitarianism presupposes the constraint, as the primary reason for accepting the luck egalitarian position over strict egalitarianism is that the latter allows for moral tyranny in a way that the former does not. Recall from Section 1.6 that the signature feature of luck egalitarianism is its incorporation of responsibility into an otherwise strict egalitarian theory of justice. There it was noted that strict egalitarian theories are, by definition, insensitive to responsibility, demanding an equal distribution of advantage regardless of anyone's past actions. By contrast, luck egalitarians are willing to declare certain inequalities just if and only if those inequalities correspond to some sanctionable choice for which the worse-off parties are responsible. And, as it turns out, there is good reason for moderating strict egalitarianism in this way: Absent this responsibility condition, strict egalitarianism is vulnerable to a species of *reductio* argument that renders the position implausible. Consider, for example, how a principle that demands strict equality of advantage would handle the case of a spiteful person who maliciously destroys any advantage she receives. In this case, a strict egalitarian would demand that advantage continually be reallocated to this person such that her share remains as great as everyone else's – a demand that is sustained even as she destroys each bit of advantage that is transferred to her until, eventually, no one has any advantage left to transfer. This result is an apparent *reductio* of strict egalitarianism.

Other examples popularly cited by luck egalitarians can be substituted into the *reductio* to reach the same conclusion. For example, consider Cohen's Aesopian case of the ant who assiduously works all summer storing up food while a neighboring grasshopper lounges idly – a decision the grasshopper makes even while recognizing that she will end up worse off when winter comes (2008, 27–8). In this case, a strict egalitarian

²¹ If the foregoing argument is correct and both the consent theory of legitimacy and the Lockean proviso follow from the moral tyranny constraint, this would reveal a little-discussed coherence in Locke's *Second Treatise on Civil Government*, which endorses both positions. Similar remarks apply to Nozick's *Anarchy, State, and Utopia*, as he defends both the Lockean proviso (1974, 178) and, in the first section of the book, something proximate to a consent theory of legitimacy, though he is not entirely consistent on this point (see Simmons 2005, 334–6).

principle of justice would demand that the ant redistribute some of her food to the grasshopper; however, there is something seemingly unfair about the ant having to make do with less because of the grasshopper's choice not to work. Thus, insofar as justice is supposed to track fairness, a strict egalitarian principle must be rejected, as it declares an unfair outcome just.²²

Why would redistribution in these cases be unfair? Cohen has suggested that such redistribution is a form of exploitation (e.g., of the assiduous ant), where exploitation runs contrary to egalitarianism (Cohen 2011, 8). However, as Michael Otsuka notes, it is unclear in what respect such exploitation can be *inegalitarian* given that the redistribution is equalizing by definition (Otsuka 2010, 223). Rather, Otsuka plausibly argues that what is unfair about such redistribution is that it forces some people to “pick up the tab” for the poor choices of others (2010, 229).

While Otsuka does not provide any analysis of what it means for someone to “pick up the tab” for someone else's choices, a natural way of specifying this notion is in terms of the moral tyranny constraint: A person has to pick up the tab for another's actions when the latter foreseeably, discretionarily, and unilaterally leaves her worse off_{FC}. Indeed, this specification explains why strict egalitarianism's prescriptions in the spiteful destroyer and lazy grasshopper cases seem intuitively unacceptable. In both cases, a strict egalitarian theory of justice licenses one party to leave others with less advantage in the counterfactual world where everyone complies with the demands of morality. Note that the spiteful destroyer and grasshopper might not, as a matter of empirical fact, leave others with less, as those others might refuse to transfer any of their advantage-producing resources. Rather, both the spiteful destroyer and the grasshopper strip others of a claim to advantage according to the strict egalitarian theory – while simultaneously acquiring a claim to that stripped advantage – such that others *would* transfer resources to these parties if the former were to fully comply with the latter's claims. Thus, the worry that motivates luck egalitarianism is that there is something problematic about a theory that allows people to leave others worse off_{FC}.

As noted in Section 1.6, the luck egalitarian solution to this problem is to *hold people responsible* for making *sanctionable choices*, where a theory

²² For similar motivating cases, see Kymlicka (2002, 73) and Stemplowska (2009, 241, 252–3). Note that such cases seem to count against Segall's (2016) thesis that there is nothing morally objectionable about any equal state of affairs. Contra Segall, the luck egalitarian theory posited by this paper will entail that equal distributions generated by sanctionable choice are unjust. Any theory that does not have this implication will problematically allow for moral tyranny.

holds someone responsible for a choice if and only if it maintains that she forfeits a claim to some quantity of advantage in virtue of that choice. By holding people responsible for their choices, luck egalitarianism is able to avoid granting the spiteful destroyer (or grasshopper) the Hohfeldian power to discretionarily, unilaterally, and foreseeably acquire a claim to others' holdings, thereby leaving them with less_{FC} . According to strict egalitarianism, when the spiteful destroyer diminishes her own advantage, she acquires a claim against others that they make equalizing transfers to her. By contrast, luck egalitarianism treats her act of destruction as a sanctionable choice in virtue of which she forfeits a claim to advantage. Specifically, there is some quantity of advantage A such that the conjunction of the destroyer having a claim to A and her act of destruction entails that others are obliged to give her some of their holdings. Luck egalitarianism holds that the destroyer forfeits her claim to A in virtue of her act of destruction, thereby precluding her from acquiring a claim to others making equalizing transfers. This, in turn, implies that fully compliant people would not make any transfers to the spiteful destroyer and, thus, no one else ends up with less_{FC} as a result of her actions. Luck egalitarianism thereby satisfies the moral tyranny constraint in a way that strict egalitarianism does not – which is to say that the constraint entails that an egalitarian theory must include a responsibility component.²³

Note that the foregoing discussion does not demonstrate that the moral tyranny constraint entails luck egalitarianism. Rather, it entails that egalitarians must endorse the luck egalitarian incorporation of responsibility, that is, the proposition that inequality is just if the worse-off party has chosen sanctionably (more on this in Chapter 6). What it does not entail is the luck egalitarian presumption that persons are entitled to equal shares of advantage absent such sanctionable choice. In other words, unlike the consent theory of legitimacy and the Lockean proviso – each of which merely constrains the kinds of claims and correlative obligations that persons can establish given any arbitrary initial set of claims/obligations – luck egalitarianism also asserts that persons start out with a claim to an equal share of advantage. Chapter 5 will argue that this thesis also

²³ This conclusion helps to elucidate Cohen's famous assertion that luck egalitarianism incorporates "within it the most powerful idea in the arsenal of the anti-egalitarian Right: the idea of choice and responsibility" (2011, 32). Specifically, the moral tyranny constraint is the foundational principle of the "anti-egalitarian Right" as it both entails core theses endorsed by right-libertarians and entails the unacceptability of strict egalitarianism. Luck egalitarianism then incorporates the constraint by ensuring that it is satisfied while still articulating a highly demanding form of egalitarianism (i.e., a position that entails that a large portion of existing inequality is unjust).

ultimately follows from a libertarian *desideratum* for moral theories (albeit not the moral tyranny constraint). For now, though, the argument will pause at the conclusion that the responsibility component of luck egalitarianism follows from the moral tyranny constraint. Such a result is seemingly sufficient for establishing the point that the constraint entails a number of influential and attractive philosophical theses and, thus, ought to be included as part of one's ultimate reflective equilibrium.²⁴

2.5 Three Objections to the Constraint

Having discussed some reasons for accepting the moral tyranny constraint, it is worth addressing three objections that might be raised against this meta-principle. The *stringency objection* posits that the moral tyranny constraint is implausible because the constraint entails that persons have a number of excessively demanding duties. For example, suppose that *P* would have married *Q* but for the fact that *Q* fell in love with *R* and got married to *R* instead. In this case, *Q*'s choice to marry *R* seemingly leaves *P* worse off_{FC} than she would have been otherwise. Thus, according to the moral tyranny constraint, a moral theory is acceptable only if it does not license *Q*'s choice to marry *R* rather than *P*, where this might be taken to require assigning *Q* a duty to marry *P*. However, given that no acceptable theory will restrict *Q*'s moral freedom in this way, a contradiction is reached. Similarly, *T* might leave *S* worse off_{FC} by opening a rival business that drives down *S*'s profits. Here, again, one might worry that the moral tyranny constraint unacceptably entails that *T* has a duty to refrain from competing with *S*. If the constraint restricts persons' moral freedom in this way, it must seemingly be rejected despite its virtues (as described in Section 2.3).

Alternatively, one might raise a *laxity objection* against the constraint. According to this objection, the constraint entails that persons do not acquire duties in cases where an adequate moral theory would assign them such duties. Consider, for example, the case where *A* and *B* are standing on

²⁴ An additional advantage of the moral tyranny constraint is that it precludes certain varieties of moral blackmail. For example, Johan E. Gustafsson (2022) notes that a committed act utilitarian (or rule utilitarian) can be successfully extorted if other agents commit to bringing about a non-utility-maximizing outcome if and only if the utilitarian does not give them money. As Gustafsson notes, "A plausible moral theory shouldn't lay one open to that kind of exploitation" (2022, 388) and the moral tyranny constraint formalizes this contention: Utilitarianism is defective because a moral theory should not allow would-be extorters to unilaterally, foreseeably, and discretionarily leave others worse off under conditions of full compliance.

the shore of a pond where a child is drowning. Assume that *B* (and not *A*) has an obligation to rescue the child, as *B* is better positioned to carry out the rescue and could do so costlessly while *A* would incur a modest cost if she were to rescue the child. Further, suppose that *B* refuses to assist the drowning child. Given *B*'s refusal to assist, it seems that *A* acquires a remedial duty to rescue the child. However, the acquisition of such a duty would violate the moral tyranny constraint, as *B* would have thereby left *A* worse off_{FC} with her choice not to rescue. Thus, the moral tyranny constraint appears to deliver the wrong results in this case.

The reply to the first objection begins with the observation that, for any given action ϕ , the moral tyranny constraint does not imply that a moral theory must declare ϕ -ing to be either permissible or impermissible. Rather, it implies that the moral theory's *entire set of posited claims* must adjust in response to an agent ϕ -ing such that no other person ends up worse off_{FC} than she would have been had the agent not ϕ -ed. Thus, the constraint does not entail that *Q* has a duty to marry *P*; rather, it entails that any loss_{FC} of advantage that *P* incurs in virtue of not marrying *Q* must be offset by some other advantage-conferring_{FC} rights assigned by the theory (where *P* would not be assigned these rights if she married *Q*). By assigning rights in this way, the posited theory of duties avoids moral tyranny by precluding *Q* from leaving *P* worse off_{FC} via her choice to marry *R* – and, crucially, it achieves this without assigning *Q* a duty to marry *P*. Similarly, the constraint does not entail that *T* has a duty to refrain from competing with *S*'s business; rather, it mandates that the theory of duties in question must preclude *T* from disadvantaging_{FC} *S* by assigning *S* other compensatory claims that, if respected, would offset any loss of advantage imposed by *T*'s choice. Thus, the stringency objection does not succeed, as the moral tyranny constraint does not entail the posited unacceptable duties.

The laxity objection poses a greater threat to the moral tyranny constraint, as it begins with the recognition that the constraint sets limits on which duties and permissions persons can have conditional on the choices that agents make. Specifically, it contends that the constraint entails an improper restriction on which remedial duties can obtain when *B* fails to discharge her duty to rescue a drowning child (namely, that *A* cannot acquire a duty to rescue the child, as such a rescue would leave *A* worse off_{FC} relative to the world where *B* chose differently). However, this objection incorrectly assumes that one must hold all other permissions and duties constant when assessing whether *A* acquires a duty to rescue the child. Were this the case, then the constraint would, indeed, imply that *A*

cannot acquire such a duty. Fortunately for proponents of the constraint, there are many alternative patterns of remedial duties that both satisfy the constraint and assign *A* a duty to rescue the child. For example, a moral theory might maintain that *A* acquires a duty to rescue the child *and* that *B* acquires a duty to compensate *A* such that *A* ends up no worse off_{FC} in virtue of *B*'s choice not to rescue the child.²⁵ In fact, one might take this conjunction of duties to be a more plausible result than simply holding that *A* has to rescue the child and must shoulder the associated costs without any compensation from *B*. If so, this putative counterexample to the moral tyranny constraint is transformed into additional reason for favoring the constraint, as the constraint can explain why it is that *B* must compensate *A* (despite the fact that her undischarged duty was owed to the *child*). Thus, the stringency and laxity objections fail to undermine the plausibility of the moral tyranny constraint.

Finally, the *paternalism objection* contends that the moral tyranny constraint is incompatible with any claim against paternalistic interference (including so-called hard paternalism that is explicitly unwanted by the beneficiary).²⁶ While not everyone believes that people have claims against paternalistic interference, such claims are widely endorsed by anarchists, libertarians, and liberals of all stripes – that is, those who would be most naturally attracted to the position advanced by the book. Thus, if the

²⁵ What if *B* is unable to compensate *A*? If one thinks that a person can still have a duty even if she is unable to discharge that duty, then *B*'s inability does not pose any special problem for the compensation solution proposed in this section. However, given that the constraint is concerned with the world of full compliance, it may well presuppose that persons can have a duty only if they are able to discharge it. If duties imply "can" in this way, then *B*'s inability to compensate *A* makes moral tyranny *inevitable*: either one assigns *A* a duty to rescue the child thereby enabling *B* to leave *A* worse off_{FC} or one does not assign *A* this duty thereby enabling *B* to leave the *child* worse off_{FC}. Given this predicament, the suggestion here is that the moral tyranny constraint should be understood to declare a moral theory unacceptable only if it violates the constraint *and* there is some rival theory that does not violate the constraint. In other words, if all possible theories entail that a person has the ability to foreseeably, discretionarily, and unilaterally leave someone worse off_{FC} by making a particular choice, then no theory should be taken to violate the constraint in virtue of that person's ability. Suppose, for example, that a nuclear-weapon-possessing villain has the ability to destroy the planet. If she makes this choice, everyone will be left worse off and there will be no way for her to compensate them. This, in turn, implies that no moral theory will be able to assign duties in such a way as to preclude her from leaving others worse off_{FC}. And, given that *no* theory can satisfy the constraint vis-à-vis this choice, it seems like the fact that some particular theory does not satisfy the constraint vis-à-vis that choice does not count against the theory. For this reason, the moral tyranny constraint should be understood to be satisfied in both this case and the aforementioned rescue case where no compensation is possible.

²⁶ In fact, one might even think that it entails that agents are *obliged* to paternalistically interfere with others, as this is the only way for a theory to avoid licensing those agents to leave others worse off_{FC}.

moral tyranny constraint implies that paternalism is acceptable, that would represent a serious theoretical cost.²⁷

The quick reply is that this objection mistakenly presupposes that persons must be assigned the set of duties such that each person ends up with the maximum possible quantity of advantage conditional on full compliance with those duties. However, the constraint merely insists that persons must *lack the ability to diminish others' advantage* conditional on full compliance. While a theory that assigned a person a claim against paternalist interference would fail to maximize her advantage_{FC} relative to an otherwise-identical theory that did not assign that claim, it does not give others any greater ability to *choose* how much advantage she ends up with under conditions of full compliance. Thus, the moral tyranny constraint does not entail that persons lack a claim against paternalistic interference.

This reply is “quick” because it does not adequately address an interesting complication that arises when the moral tyranny constraint is applied to theories that include claims against such interference. Consider an arbitrary moral theory that assigns Q a claim against P ϕ -ing where ϕ -ing is an act of paternalistic interference. Because ϕ -ing is an act of paternalism, P ϕ -ing will leave Q better off than if she discharges her duty and does not ϕ – which is to say that she leaves Q better off in the actual world A than in the world of full compliance F_1 . However, her choice to ϕ also changes what the full-compliance world looks like because P will acquire remedial duties in virtue of her failure to discharge her duty to not paternalistically interfere with Q (call this adjusted full compliance world F_2). Specifically, one might think that P will acquire a new duty to *compensate* Q for failing to discharge her duty to Q , where this compensatory remedial duty entails that Q is better off in F_2 than she is in A . Further, given that Q is better off in A than she is in F_1 , transitivity implies that Q is better off in F_2 than she is in F_1 . Thus, Q is better off_{FC} if P ϕ -s than if P discharges her duty and does not ϕ . This result may seem like a problem for the moral tyranny constraint because it appears that the moral theory – which is to say, any moral theory that includes a claim against paternalistic interference – violates the constraint, as P can unilaterally leave Q worse off_{FC} by doing her duty and refraining from ϕ -ing. One might therefore conclude that the constraint does, in fact, problematically imply that there are no claims against paternalistic interference.

Against this worry, note that the moral tyranny constraint does not merely hold that a theory cannot license a person to leave others worse

²⁷ For a recent defense of paternalism, see Jason Hanna (2018).

off_{FC}. Rather, it maintains that the theory must not license her to *unilaterally, foreseeably, and discretionarily* leave others worse off_{FC}. For these purposes, the last qualifier is the crucial one, as *P* discharging her duty to refrain from ϕ -ing would not qualify as discretionary in the sense described in Section 2.2. There it was stipulated that a person acts discretionarily only if she does not *have* to carry out that action, where “have” can be interpreted either in terms of ability – that is, there is no other option physically available to the agent – or normatively, which is to say that the agent lacks a permission to carry out any rival action. Because *P*'s duty to refrain from ϕ -ing entails that she lacks a permission to ϕ , it follows that she does not discretionarily leave *Q* worse off_{FC} by declining to ϕ (although she does, in fact, leave *Q* worse off_{FC}). Thus, the posited theory does not violate the moral tyranny constraint, which, in turn, implies the more general conclusion that the constraint is compatible with duties against paternalistic interference.

2.6 Conclusion

This chapter has taken the first steps toward demonstrating the coherence of social anarchism. Specifically, it has argued that three of the position's theses follow, either in part or in whole, from a more general constraint on which normative theories are acceptable, namely, the moral tyranny constraint. The chapter has thereby demonstrated that these theses are neither incompatible nor an arbitrary set of views conjoined together without reason. Rather, they are logically connected in a way that renders the social anarchist position coherent in the sense defined in Section I.2. Additionally, this chapter has argued that the moral tyranny constraint is independently plausible. Given its plausibility, the fact that the constraint entails consent theory, the Lockean proviso, and luck egalitarianism's incorporation of responsibility is a reason to accept these positions. The task of the next three chapters is to show that the remaining anarchist theses – namely, the self-ownership thesis, the anarchist rejection of private property, and the anarchist conclusion – similarly cohere with the other components of the anarchist position.