

Luck Egalitarianism without Moral Tyranny

We want to put everything in common, starting from the principle that everybody should do some work and all should live as well as possible. It's not possible to live in this world without working, so if one person doesn't do anything he has to live at the expense of others, which is unfair and harmful. Obviously when I say that everybody should work I mean all those that are able to, and do the amount suited to them. The [disabled], the weak and the aged should be supported by society, because it is the duty of humanity that no one should suffer.

Errico Malatesta, *Between Peasants*

Chapters 3, 4, and 5 have attempted to provide a libertarian argument for luck egalitarianism.¹ Chapters 3 and 4 each took a prominent libertarian thesis and argued that it entails that there are no existing private property rights. Chapter 5 then argued that, in the absence of property, libertarians ought to embrace the anarchist conclusion, which assigns persons luck egalitarian distributive claims over unowned natural resources. More precisely, this conclusion posits that each person has a claim against others interacting with unowned resources in a way that (a) would leave her worse off than someone else where (b) this comparative disadvantage does not appropriately correspond to previous sanctionable choices on her part. The obvious – and so far unanswered – question is: What choices count as sanctionable and which inequalities can be said to appropriately correspond to those choices? This chapter seeks to answer both parts of this question by appealing to the moral tyranny constraint.

Recall from Section 2.4 that luck egalitarianism avoids the moral tyranny of strict egalitarianism by holding people responsible for making sanctionable

¹ This chapter is an adapted version of a paper originally published in *Philosophical Studies* (Spafford 2022). While the animating idea is the same in both versions, some of the technical details of the paper have been adjusted here to improve the proposed theory (particularly in Sections 3–5).

choices. The problem with strict egalitarianism is that it allows a spiteful destroyer to preserve her claim to an equal share of advantage even as she destroys any advantage in her possession. Because the destroyer preserves this claim, full compliance would require that others make costly equalizing transfers to her. She is thereby able to discretionarily, foreseeably, and unilaterally leave those others worse off_{FC} when she destroys her holdings. By contrast, luck egalitarianism declares that the destruction of her holdings is a sanctionable choice, which, in turn, implies that she forfeits a claim to some portion of advantage (where this forfeiture negates her claim to any transfer). The theory thereby holds her responsible for her sanctionable choice and, in this way, satisfies the moral tyranny constraint.

However, as was noted in Section 1.6, there are two respects in which this general idea of holding people responsible needs to be specified if the anarchist conclusion is to have determinate content. First: Which choices qualify as sanctionable? Some general theory is needed here that can be applied to cases like the spiteful destroyer or Cohen's parable of the ant and the grasshopper. Second: To what quantity of advantage does a person lose her claim when she makes a sanctionable choice? Most luck egalitarians answer these questions by either explicitly or tacitly assuming what might be called *prudential contextualism*: A person chooses sanctionably if and only if she is responsible for leaving herself worse off than she could have been otherwise.² She then forfeits a claim to however much additional advantage she would have had if she had chosen differently. For example, if the norm in a given society is that people who park on a certain street have their car towed, then the person who parks there and has her car towed forfeits her claim to the extra advantage she would have possessed had her car not been towed.

Unfortunately, this interpretation of luck egalitarianism has left the position vulnerable to three serious objections – each of which would equally apply to the anarchist conclusion. Most notably, one of these objections holds that this interpretation puts both luck egalitarianism

² The term “prudential contextualism” is a slight modification of Olsaretti's term “contextualism,” which she uses to describe this dominant interpretation of luck egalitarianism (2009, 180). The chapter will not take a stand on the exact sense in which an agent must be *responsible* for leaving herself worse off – at least, beyond a few claims about the necessary conditions of such responsibility discussed subsequently. It should be noted that not all luck egalitarians are prudential contextualists with some explicitly rejecting the position (e.g., Olsaretti (2009), Stemplowska (2009), and Thaysen and Albertsen (2017)). These alternative views will also be discussed. It should also be noted that there is some dispute over whether certain influential luck egalitarians were contextualists. For a critical discussion, see Stemplowska (2013).

and the associated anarchist conclusion in violation of the moral tyranny constraint. The task of this chapter is to present an alternative account of sanctionable choice that avoids these objections. Specifically, it will propose that an agent chooses sanctionably if and only if the choice, under conditions of full compliance, can reasonably be expected to produce less appropriately distributed advantage than some alternative choice that could have been made. The remainder of the chapter will explain what is meant by each of these terms and how the proposed account resolves the three objections. First, though, these objections need to be introduced, with particular attention paid to the moral tyranny objection.

6.1 Three Objections to Prudential Contextualism

The first objection to contextualist luck egalitarianism has been forcefully raised by Richard Arneson, who argues that luck egalitarianism delivers incorrect results in cases of charitable action. For example, a Mother Teresa figure who impoverishes herself assisting the poor would, on the prudential contextualist view, have made a sanctionable choice, as she leaves herself worse off in a way that could have been avoided. Thus, the luck egalitarian (or social anarchist) who endorses prudential contextualism would seemingly be committed to saying that she is not entitled to any sort of compensatory redistribution; however, Arneson contends that this is the wrong result, with luck egalitarianism then being rejected as part of a *modus tollens* argument (Arneson 2011a, 244; 2011b, 33–4).³

Second, there is what Susan Hurley has called luck egalitarianism's "boring problem." This objection aims to call into question the core luck egalitarian contention that sanctionable choice can justify inequality. Hurley argues that this contention is plausible only if sanctionable choosers are responsible for the inequality in question. However, on the prudential contextualist view, a person chooses sanctionably if and only if she is responsible for *her particular level of advantage*. Given that someone can be responsible for her level of advantage but not the associated inequality – as the inequality is partly a function of the advantage levels of others for which *they* are responsible – it follows that sanctionable choice cannot justify inequality in the way that luck egalitarians contend (Hurley

³ For an alternative version of this argument, see Larry Temkin's case of a good Samaritan who rescues a drowning child from a pond but injures herself in the process (2011, 63). Here, again, it is maintained that her failure to act prudently means she is held responsible for this personal cost via the denial of any compensation.

2003, 160–1).⁴ Similarly, insofar as the anarchist is motivated by egalitarian concerns, she will want to posit that, while there is something *prima facie* unjust about inequality, there is nothing unjust about inequality corresponding to sanctionable choice. Thus, her position will similarly be vulnerable to the objection that sanctionable choice fails to defeat the *prima facie* injustice of inequality.

Finally – and most importantly for the purposes of this chapter – there is Serena Olsaretti’s objection (though she does not use this language) that contextualist luck egalitarianism violates the moral tyranny constraint. Specifically, she begins with the observation that most luck egalitarians tacitly assume a contextualist theory of sanctionable choice where a person forfeits a claim to however much advantage she foregoes due to contingent social circumstances and the choices that others make (2009, 180). However, given such a prudential contextualist view, a person’s claim to advantage will often – and problematically – depend on the capricious choices of others (176). To illustrate this point, Olsaretti introduces Marc Fleurbaey’s (1995) case of a reckless motorcyclist who crashes and is injured as a result of driving too fast and not wearing a helmet. She notes that, on the contextualist theory, the quantity of advantage to which the motorcyclist forfeits a claim will be a function of whether a passerby provides assistance, leaves her unaided, or confiscates her motorcycle – a result that makes contextualist luck egalitarianism seem unacceptable (2009, 175–6).

While Olsaretti does not provide a general theory of why this is a problem for contextualism, one can provide such an explanation by appealing to the moral tyranny constraint. Specifically, contextualism allows the passerby to unilaterally, foreseeably, and discretionarily reduce the quantity of advantage to which the motorcyclist has a claim which, in turn, would reduce how much advantage she would have in the world of full compliance. If the passerby declines to assist the motorcyclist, the latter will be left with less_{FC} advantage than if she would possess if she were assisted. And she would be left with less_{FC} still if the passerby were to confiscate her motorcycle. Contextualist luck egalitarianism thereby violates the moral tyranny constraint, as it enables the passerby to unilaterally, discretionarily, and foreseeably leave the motorcyclist worse off_{FC}. This result explains why prudential contextualism is an unacceptable theory of

⁴ For a recent paper that attempts to extend the boring problem into a more general objection to theories that declare luck-based inequality unjust, see Matthew T. Jeffers (2020). For an alternative reply to the boring problem, see Spafford (2023).

which choices count as sanctionable (particularly given the fact that satisfying the constraint was the motivating reason for adopting luck egalitarianism in the first place). Thus, the anarchist conclusion must employ some alternative account of sanctionable choice – ideally, one that also resolves the other two objections to luck egalitarianism discussed just prior.

6.2 Moralized Contextualism

What is needed is an alternative account of sanctionable choice that satisfies the moral tyranny constraint. The task of the remainder of the chapter is to provide such an account. First, though, it is worth briefly considering an alternative account that might seem like a simple solution to contextualism's moral tyranny problem. This account posits that a person chooses sanctionably if and only if (a) she leaves herself worse off than she would have been otherwise and (b) *she does not end up worse off as a result of someone infringing upon her rights*. She then forfeits a claim to the surplus advantage that she would have possessed had she chosen differently.⁵

This *moralized contextualism* seems to rule out some of the more intuitively problematic forms of moral tyranny endorsed by standard contextualist accounts. For example, in the motorcycle case, one might think that Condition (b) is not met, as the passerby who confiscates the motorcycle infringes upon the rights of the motorcyclist to continue to use her motorcycle. Thus, the motorcyclist does not choose sanctionably, which, in turn, implies that she does not forfeit a claim to the advantage that would result from her continued use of the motorcycle. Further, given that she still has a claim to this advantage, full compliance with her claims would ensure that she retains possession of this advantage (e.g., because the passerby would immediately return the motorcycle), thereby preventing the passerby from leaving the motorcyclist with less_{FC}. This result suggests that moralized contextualism satisfies the moral tyranny constraint.

⁵ This seems to be what Olsaretti is suggesting when she says that “the notion of responsibility a theory of justice employs is necessarily moralized, in that it must presuppose a view of what individuals owe to one another in order to determine the legitimate consequences of choices” (2009, 186). A more formal articulation of moralized contextualism is proposed by Zofia Stemplowska (2009), though she significantly qualifies the view. Unfortunately, working through the interesting details of her view would take things too far afield.

However, there are two problems with this suggestion. First, even if moralization *limits* the extent to which a person is able to leave others with less_{FC}, it does not *eliminate* her ability to leave them with less_{FC}. For example, while the motorcyclist may have a right to her motorcycle, she does not obviously have a right to assistance after the accident (at least, if assisting would be reasonably costly to the passerby). Given the absence of such a right, moralized contextualism delivers the same result as standard contextualism in the case where the passerby chooses not to assist the motorcyclist: The motorcyclist forfeits a claim to however much advantage she fails to secure as a result of the passerby's choice. Thus, full compliance under moralized contextualism would still allow the passerby to (unilaterally, foreseeably, and discretionarily) leave the motorcyclist with less_{FC}. The fact that moralized contextualism grants the passerby this ability entails that the theory still violates the moral tyranny constraint.

The second problem with the moralized contextualist approach is that it would render the anarchist conclusion circular. Note that the anarchist conclusion is supposed to answer the question of which rights people have over objects and resources. Indeed, this is the very point at issue in debates over taxes and transfers, with anarchists contending that natural resources should be distributed in a responsibility-sensitive egalitarian fashion. However, one cannot then assume that there is a given set of rights over objects for the anarchist to use as an input for her theory. For example, it cannot be maintained that the passerby infringes upon the motorcyclist's right to use her motorcycle, as it is an open question whether the motorcyclist does, in fact, have a right to use that motorcycle (particularly given her choice to ride without a helmet). Thus, an anarchist position that assigns claims in accordance with moralized contextualist luck egalitarianism not only fails to satisfy the moral tyranny constraint, but also appears to be unacceptably circular.

6.3 A Theory of Sanctionable Choice

Given the unacceptability of (moralized) contextualism, anarchists need a theory of sanctionable choice that satisfies the moral tyranny constraint. This section proposes the following account: An agent chooses sanctionably if and only if she fails to maximize *warranted expected distributed advantage assuming full compliance*. The task of this section is to explicate each of these italicized concepts and explain why the theory defines sanctionable choice in this way.

To begin, recall that a theory satisfies the moral tyranny constraint if and only if it does not enable any person to unilaterally, foreseeably, and discretionarily leave others with less_{FC}. Further, recall that there are two ways that a theory might enable a person in this way. First, it might grant her the power to directly strip others of their claims to advantage (or impose advantage-diminishing obligations on them), thereby changing how much advantage they would possess under conditions of full compliance. Second – and more importantly for these purposes – it might fail to adequately sanction those who diminish the total quantity of advantage that would be available if all persons were to fully comply with the demands of morality. To put this point a bit more precisely, such a failure occurs when a person reduces_{FC} the total quantity of advantage by some quantity x but the theory in question holds that her just share is diminished by a quantity that is less than x . Indeed, this is what a strict egalitarian theory asserts in the case of the spiteful destroyer: Even as the destroyer reduces_{FC} the total amount of available advantage by x , the theory holds that she only forfeits a claim to a quantity of x/n , where n is the number of people in the scenario. As a result, if others were to respect her adjusted claim to advantage, at least some persons would end up with less advantage than they would have had otherwise.

To avoid moral tyranny, then, a theory must hold that those who reduce_{FC} the total quantity of advantage forfeit a claim to a sufficient quantity of advantage such that compliance with their claims would not leave others worse off. In other words, when people reduce_{FC} the total quantity of advantage, the theory must declare that they choose sanctionably and hold them responsible by reducing the quantity of advantage to which they are entitled, thereby making them internalize the costs_{FC} of their actions. This is the core idea of the theory of sanctionable choice presented here. However, some additional groundwork and a few refinements are needed to both make the theory acceptably egalitarian and ensure that the correct people are held responsible.

To further explicate the theory, it will be helpful to stake out a position regarding the *equilisandum* of the anarchist conclusion (i.e., the thing to which persons have equal distributive claims). Specifically, it will be assumed that what is to be equalized is *lifetime* levels of advantage, as opposed to the advantage persons possess at a particular time or over some specified period. Thus, there is no injustice in an arrangement where one person labors for the first half of her life while another relaxes, so long as the two switch roles for the second halves of their respective lives. Early on,

the two people will have very different quantities of advantage; however, injustice only obtains if things are not adequately evened out in the future.⁶ Given this assumption, any future use of terms like “the distribution of advantage” should be understood as referring to how lifetime advantage is distributed.

With this simplifying assumption in place, it becomes possible to determine the quantity of destroyed_{FC} advantage that an agent must internalize. A natural temptation is to simply calculate the total quantity of advantage destroyed_{FC} by her choice by taking the total advantage_{FC} that obtains given her choice and comparing it to the maximum quantity that would have obtained had she chosen differently (assuming full compliance in both cases). However, this approach is unacceptable for two reasons. First, it would still violate the moral tyranny constraint. Note that a consequence of using this method is that later choices by other parties will sometimes determine how much destroyed_{FC} advantage an agent has to internalize. Thus, a theory that calculates lost advantage in this way grants later choosers the power to unilaterally, discretionarily, and foreseeably render an agent’s earlier choice sanctionable, thereby leaving her with less_{FC}. Second, this approach would make sanctionable choice a function of luck, as a choice might reduce_{FC} the quantity of total advantage more than a rival choice due to an unforeseeable future event. Assuming that sanctionable choice requires that the agent be responsible for the state of affairs that grounds the forfeiture of her claim to advantage – and given the fact that agents are not responsible for the unforeseeable consequences of their actions – it follows that the sanctionability of a choice cannot be a function of whether or not it reduces_{FC} the total quantity of advantage relative to a counterfactual choice.⁷

⁶ This assumption helps to simplify things in the following way. Later, there will be much talk of how advantage is distributed. If the *equilibrandum* of the luck egalitarian principle is lifetime advantage, then there is only one distribution to be assessed, namely, the lifetime levels of advantage everyone ends up with. By contrast, alternative approaches entail that there are many distributions that obtain across time, each of which would have to be assessed. That said, there are objections to taking entire lives as the basic unit of egalitarianism. See, for example, McKerlie (1989) and Temkin (1993).

⁷ This assumption is posited because forfeiture seems unacceptably arbitrary if it is not grounded in some kind of responsible choice. Absent responsibility as a necessary condition, it seems one might equally forfeit claims to advantage in virtue of others’ choices. Granted, such a condition does allow persons to act in ways that leave others with less_{FC}, as it allows them to sustain their claim to having as much advantage as everyone else despite diminishing the total quantity of available_{FC} advantage. However, it does not allow them to *foreseeably* leave others with less_{FC}, thereby avoiding any contradiction with the moral tyranny constraint.

The observation that responsibility requires foreseeability suggests an alternative approach to quantifying how much destroyed_{FC} advantage an agent must internalize: Instead of determining whether her choice reduces_{FC} the total quantity of available advantage, one must determine whether that choice is *expected* to reduce_{FC} the total quantity of advantage. Note that, at any point in time, there is a set of possible ultimate distributions of advantage that could still arise given all of the preceding events (where distributions extend across lifetimes as described earlier). Further, for any given choice that an agent might make, each of those distributions will have a particular probability of obtaining conditional on that choice being made and future full compliance. These conditional probabilities make it possible to calculate the *expected advantage* of that choice (assuming full compliance) by taking the total quantity of advantage of each distribution, multiplying it by the aforementioned conditional probability, and summing the results. This value can then be compared to the expected advantage value of other rival choices that could have been made. When a person makes a choice that has a lower expected advantage value than a rival choice that she could have made, she can be preliminarily understood as having chosen sanctionably, as her action is expected to diminish_{FC} the total quantity of advantage.⁸

This account of sanctionable choice is merely preliminary because a significant revision must be made vis-à-vis calculations of expected value. To calculate the expected value of a choice, one must assign an *advantage value* to each of the possible distributions and then multiply that value by the probability of that distribution obtaining conditional on the choice being made and future full compliance. So far, this advantage value has been set equal to the total quantity of advantage that obtains in that distribution (as bringing about a distribution with less total advantage leaves others with less_{FC}, *ceteris paribus*). However, using total advantage obscures how advantage is *distributed* across persons. This is a problem because the moral tyranny constraint requires that agents must not be able to leave *any person* with less_{FC} than she would have had otherwise, not *people on average*. But agents *will* be able to leave particular persons with less if sanctionable choices are specified to be all and only those choices that do not maximize expected value.

⁸ This use of expected value has been embraced by a number of luck egalitarians including Arneson (1989), Knight (2013), and Vallentyne (2002; 2008).

To see this, consider the case where agent P can either ϕ or ψ . If she ψ -s, she will realize a distribution where she, Q , and R each end up with 10 units of advantage. Alternatively, if she ϕ -s, she will realize a state of affairs where R has 20 units of nontransferrable advantage and she and Q have 16 units to split between the two of them (at P 's discretion). In this scenario, there are two possibilities: either P 's ϕ -ing is sanctionable or it is not. If P 's ϕ -ing does not count as sanctionable, then P would retain her claim to an equal share of the available distributable advantage (8 units).⁹ Thus, in the world of full compliance, she keeps 8 units for herself and similarly leaves Q with 8 units of advantage – that is, P is able to leave Q worse off_{FC} by ϕ -ing relative to the world where P had chosen to ψ instead. By contrast, if P 's choice to ϕ is sanctionable, then she loses her claim to a full 8 units of advantage, thereby allowing a portion of that advantage to be reassigned to Q such that Q would receive 10 units of advantage under conditions of full compliance. Given this result, it follows that P 's choice to ϕ must be declared sanctionable if the proposed theory is to satisfy the moral tyranny constraint. However, note that P ϕ -ing also maximizes the expected total quantity of advantage (by producing 36 units of advantage rather than the 30 produced by ψ -ing). Thus, an acceptable theory cannot hold that a person chooses sanctionably if and only if she fails to maximize_{FC} expected total advantage.¹⁰

⁹ It is assumed here that P and Q have an equal claim to the quantity of *distributable* advantage; i.e., they each have a claim to 8 units of the 16 that can be split between them. Alternatively, one might maintain that P and Q have a claim to an equal share of the *total* advantage, which is to say they each have a claim to 12 units of advantage. However, the former approach is endorsed here because it simplifies some of the subsequent discussion and also avoids any impossibility of rights (i.e., cases where two rights cannot be simultaneously respected). Either way, the following point remains true about this case: P is able to leave Q with less_{FC} by ϕ -ing relative to ψ -ing if her choice to ϕ is non-sanctionable.

¹⁰ It is worth noting that P ψ -ing will also leave R with less_{FC} than R would have had if P had ϕ -ed. Thus, one might worry that moral tyranny is inevitable in cases where advantage is nontransferrable. However, this concern can be sidestepped by qualifying the moral tyranny constraint such that moral tyranny does not obtain if the person who is left with less_{FC} (a) ends up with a just share of advantage and (b) is only left worse off relative to a counterfactual where full compliance would have delivered her a quantity of advantage that exceeded her just share (due to agents' limited ability to transfer). Indeed, there is seemingly nothing problematic about a theory that enables agents to deny full compliers advantage exceeding that to which they have a claim. For the sake of parsimony, though, the rest of the chapter will gloss over this qualification.

Alternatively, one might hold that while both ϕ -ing and ψ -ing leave someone with less_{FC}, P does not *discretionarily* leave R with less when she ψ -s because ψ -ing is the only way to respect Q 's claim to receiving an equal share of advantage (in addition to everyone else). Given that ϕ -ing does not respect Q 's claim in this way, P is obliged to ψ . Thus, it is only ϕ -ing that is problematic vis-à-vis the

This observation demands that the theory be refined as follows: A sanctionable choice does not merely fail to maximize the expected total advantage under conditions of full compliance but, rather, fails to maximize the total quantity of *appropriately distributed advantage* under conditions of full compliance. The idea here is as follows. At any given point in time, a luck egalitarian theory of justice – and, by extension, the anarchist assignment of distributive claims – will dictate what quantities of possessed advantage are just: Each person is entitled to as much advantage as everyone else minus however much advantage she has forfeited due to previous sanctionable choices. When each person possesses her just share, advantage can be said to be appropriately distributed, with full compliance entailing that people do everything permissible to realize this state of affairs. The problem is that certain choices preclude the possibility of fully compliant persons bringing about an appropriate distribution of advantage, with some persons thereby ending up with less_{FC} than they could have had otherwise (and less than the amount of advantage to which they have a claim). To avoid this form of moral tyranny, the theory must treat such choices as sanctionable *in addition to* choices that leave others with less_{FC} by failing to maximize $_{\text{FC}}$ the total quantity of transferrable advantage.

To deliver the result that such choices are sanctionable, one can modify the advantage values used to calculate the expected advantage of a choice. Specifically, any distribution where advantage is appropriately distributed receives an advantage value equal to the total quantity of advantage (i.e., the sum of each person's lifetime advantage). By contrast, for any distribution U where advantage is inappropriately distributed, U 's advantage value is calculated via the following procedure. First, of those distributions that have a nonzero probability conditional on all past events obtaining, identify the distribution E that has the greatest total appropriately distributed advantage. Second, identify all those persons in U who have less advantage than they would have had in E . Third, sum the differences between how much advantage each such person has in E and how much she has in U . Fourth, calculate the advantage value of U by subtracting this sum from the total value of E (as this reduction reflects how much less_{FC} advantage people end up with in U relative to the counterfactual E where

moral tyranny constraint, as the worsening $_{\text{FC}}$ due to ψ -ing is nondiscretionary given that ψ -ing is obligatory. This means that a theory of sanctionable choice should only treat P 's ϕ -ing and the associated loss_{FC} of advantage as sanctionable. This is the aim of the theoretical refinement proposed in the next paragraph. Additionally, Section 6.6 will say more about the relationship between obligatory actions and the proposed theory of sanctionable choice.

their just shares are maximized). Fifth, weight the advantage value of each distribution by multiplying it by the probability that the distribution obtains conditional on the choice in question being made and all persons complying with the demands of justice going forward.¹¹ Finally, calculate the *expected distributed advantage value* of the choice by summing those weighted advantage values.

Once the expected distributed advantage value of each possible choice has been calculated, it becomes possible to compare the value of the actual choice to those of rival choices that could have been made. When the former is less than one of latter values, the agent is responsible for leaving everyone with less_{FC}, where this difference quantifies the total advantage lost_{FC} due to her choice. More precisely, everyone's combined loss_{FC} of advantage is equal to the absolute value of the difference between the expected distributed advantage value of her choice and the value of the choice with the maximal expected distributed advantage value.

Note that the proposed theory of sanctionable choice does not need to prevent an agent from leaving *everyone* with less_{FC}. Rather, to satisfy the moral tyranny constraint, it must merely preclude her from leaving *others* with less_{FC}. Thus, the relevant question is what portion of the expected total loss_{FC} of distributed advantage would be imposed upon others if the agent were not held responsible for her choice. Fortunately, calculating this value is fairly straightforward. Because the anarchist conclusion is egalitarian in character, it holds that, in the absence of sanctionable choice, any diminution in the total stock of advantage is distributed equally across persons. Thus, if the total quantity of lost_{FC} advantage is equal to x and there are n persons in the world, each person will absorb $1/n$ of that lost_{FC} advantage, that is, $\frac{1}{n} \times x$.¹² Given that the total number of people who are not the agent – that is, all those upon whom the cost_{FC} would be

¹¹ Note that to avoid circularity, the theory must maintain that, when determining what fully compliant people will do given some choice, it must be assumed that they will treat the choice as non-sanctionable. Otherwise the theory will problematically maintain that the sanctionability of a choice depends on what fully compliant people will do in response to that choice, which, in turn, depends on whether the choice is sanctionable. This stipulation is perhaps a bit ad hoc, but this slight theoretical vice does not seem like a significant problem given the account's many significant theoretical virtues.

¹² This is a slight oversimplification, as it ignores cases where some people receive a quantity of nontransferrable advantage that is either equal to or exceeds their just share. In such cases, they will not absorb any of the cost imposed by the agent (because none of their advantage can be transferred away). Thus, rather than standing for the total number of people, " n " should really stand for the total number of people minus those whose nontransferrable advantage insulates them from the effects_{FC} of any choice-responsive adjustment of claims.

imposed – is equal to $n - 1$, the sum of all of their incurred costs_{FC} would then be equal to $\frac{n-1}{n} \times x$. This value represents the costs_{FC} that the agent would impose upon others absent any forfeiture. The theory of sanctionable choice then holds the agent responsible by asserting that she forfeits a claim to this same quantity of advantage – that is, this forfeited quantity is subtracted from the quantity of advantage to which she would have been entitled absent forfeiture. Additionally, everyone else acquires a claim to a share of the total quantity of forfeited advantage such that they are not left worse off_{FC} by the agent's choice.¹³ Together, this forfeiture and accompanying claim acquisition forces the agent to internalize_{FC} any foreseeable losses_{FC} of advantage attributable to her action (by effectively transferring her claim to this advantage to those who would otherwise be left worse off_{FC}). Thus, the proposed theory of sanctionable choice precludes the agent from foreseeably leaving others with less_{FC} and thereby satisfies the moral tyranny constraint.

One final bit of elaboration is needed to complete the formal account of sanctionable choice. So far, sanctionable choice has been defined in terms of expected distributed advantage, where this value is a function of distributions' distributed advantage values and the conditional probabilities of those distributions obtaining. However, note that the term "probability" is ambiguous. On the one hand, it might refer to *objective* probabilities, which, in this case, represent how likely it is in some metaphysical sense that a distribution will arise. Alternatively, it might refer to *subjective* probabilities representing the agent's beliefs about how probable it is that a distribution will arise. To eliminate this ambiguity, one can adopt Carl Knight's suggestion that the proper account of probability to incorporate into luck egalitarian (and, in this context, anarchist) calculations of expected value is *warranted subjective probability adjusted for non-culpable incapacity* (2013, 1067). Briefly, Knight contends that the relevant probability is that which the agent *should* have assigned given the evidence available to her – at least, in those cases where she is capable of assessing that evidence (1067). The advantage of this evidentialist view is that it does not differentially hold people responsible for their unchosen epistemic states. By contrast, on the objective probabilities approach, people might be unaware of relevant objective probabilities "through no

¹³ For the sake of concision, the chapter will only talk of the sanctionable party forfeiting a claim going forward. However, this should always be taken as shorthand for the assertion that the sanctionable party forfeits a claim *and* the relevant non-sanctionable parties acquire claims in the way just described.

fault or choice of their own,” and this makes it unfair to hold them responsible for non-maximizing choices (1066). Similarly, using subjective probabilities is unfair because those who are, as a matter of luck, overly confident that their choices are optimal will be found less culpable than those who lack that confidence (1066).

For the purposes of this argument, it is helpful to restate Knight’s argument in terms of responsibility. An agent cannot seemingly be responsible for that which she could not have known given the evidence available to her. Further, even if she could know certain things, she may not be responsible for failing to form the proper beliefs given certain extenuating circumstances.¹⁴ Given that responsibility is a necessary condition of sanctionable choice, it follows that an account of sanctionable choice that incorporates calculations of expected value should adopt Knight’s notion of probability.¹⁵ Thus, sanctionable choice should be understood as a failure to maximize *warranted* expected distributed advantage (WEDA) under conditions of full compliance.¹⁶

6.4 Applying the Theory

This abstract description of the theory can be illustrated by applying it to a highly simplified version of Fleurbaey’s (1995) reckless motorcyclist case. Specifically, it will help to provide some invented numbers to demonstrate how the relevant calculations are carried out, beginning with the WEDA value of the motorcyclist choosing to wear a helmet:

¹⁴ For more on this point, see Vallentyne (2002, 536).

¹⁵ One might slightly amend Knight’s account in the following way. Knight suggests that, in cases where an agent is not culpable for her failure to assess the evidence, she should be treated as having not made a choice at all, and, thus, not acted sanctionably (2013, 1068). However, one might alternatively think that, in cases where an agent is not responsible for her incorrect beliefs about how likely various distributions are to obtain, she might still be responsible for making a suboptimal choice *given* those beliefs. Thus, one might calculate expected distributed value using subjective probabilities in such cases.

¹⁶ One consequence of incorporating Knight’s suggestion is that the proposed theory of sanctionable choice does not countenance *option luck*. Briefly, option luck is generally understood to be the outcome that results from a deliberate and avoidable gamble, with many luck egalitarian theories taking persons to have chosen sanctionably if (a) they choose to make such a gamble and (b) they lose out on advantage as a result (see, e.g., Dworkin (1981, 293)). In other words, if a person gambles and loses, this is judged to be bad option luck for which the person is held responsible (i.e., the fact that she ends up worse off than others is held to be just by the theory). By contrast, the proposed theory does not hold people responsible for losing a gamble; rather, it holds them responsible for making any gamble that does not maximize WEDA, irrespective of how that gamble turns out. It is, thus, a variety of what Shlomi Segall has called “all-luck egalitarianism” (2010, 46).

Table 6.1 *Calculating the WEDA Value of Motorcycling with a Helmet*

	Distribution 1	Distribution 2	Distribution 3	Distribution 4
	Wears a Helmet			
	Does Not Crash	Crashes		
	No Transfer	Assist (No Transfer)	Assist (Transfer)	Failure to Assist
Distributed Advantage	M = 500, P = 500	M = 480, P = 500	M = 490, P = 490	M = 360, P = 500
Advantage Value	1000	980	980	860
$P(\text{Distribution} \mid \text{Helmet})$	0.7	0	0.1	0.2
Expected Advantage	700	0	98	172
WEDA	970			

Each column of Table 6.1 represents a distribution, where that distribution is defined in terms of the unique set of events compatible with the ultimate quantity of advantage possessed by all persons (in this case, the motorcyclist and the passerby). In Distribution 1, the motorcyclist wears her helmet, does not crash, and no subsequent transfers of advantage (or additional events) occur. In Distribution 2, the motorcyclist wears a helmet, crashes, and is then assisted by a passerby. Distribution 3 is defined by the same series of events as Distribution 2 except the passerby also makes an equalizing transfer to the motorcyclist. And, finally, in Distribution 4, the passerby simply drives past the injured motorcyclist and does not assist her. (For simplicity, assume that the passerby cannot help the motorcyclist.)¹⁷

The first row of Table 6.1 represents how much advantage each person is stipulated to possess in each distribution. In Distribution 1, where the motorcyclist does not crash, she and the passerby each end up with 500 units of advantage. In Distribution 2, the passerby is able to costlessly treat the motorcyclist's injuries from the crash and, thus, is left with the same quantity of advantage that she would have had if no crash had occurred (500 units). The motorcyclist, however, is a bit bruised and battered, so she ends up with only 480 units of advantage. In Distribution 3, the passerby's supplemental transfer increases the motorcyclist's advantage

¹⁷ Additionally, it will be assumed that the passerby is unable to make any transfer to the motorcyclist in this scenario, perhaps because she continues traveling to a distant location where her spatiotemporal position makes it impossible to relocate the motorcyclist and transfer advantage to her.

to 490, but that comes at the expense of the passerby, who also ends up with 490 units of advantage. Finally, if no assistance is given, the passerby maintains her original 500 units of advantage while the motorcyclist's untreated injuries reduce her advantage to 360 units.

The second row of Table 6.1 represents the advantage value of each of the four distributions, where this value is a function of the values listed in the first row. As noted earlier, the first step of calculating the advantage value of a distribution is to identify the distribution with the greatest total quantity of advantage that is appropriately distributed and that has a nonzero probability of obtaining conditional on the choice under consideration being made under conditions of full compliance. Assuming no prior sanctionable choice on the part of either party, Distribution 1 satisfies these conditions, and, thus, receives an advantage value equal to the total advantage possessed by all persons (1000 units). The next step is to calculate the advantage value of each additional distribution by identifying every person in that distribution who ends up worse off than she would have been in the comparison distribution. In Distributions 2 and 4, only the motorcyclist ends up worse off, while in Distribution 3, both the motorcyclist and the passerby end up worse off. Each difference in advantage is then subtracted from the advantage value of the comparison distribution. So, for Distribution 3, one would subtract 10 (the difference between how much advantage the motorcyclist has in Distribution 1 and how much she has in Distribution 3) and another 10 (the difference between how much advantage the passerby has in Distribution 1 and how much she has in Distribution 3) from 1000 to get an advantage value of 980.

The third row of Table 6.1 represents the warranted probability of each distribution obtaining conditional upon the motorcyclist not wearing a helmet under conditions of full compliance. For the purposes of this example, it is stipulated that the motorcyclist's evidence suggests there is a probability of 0.7 that Distribution 1 obtains, a probability of 0 that Distribution 2 obtains, a probability of 0.1 that Distribution 3 obtains, and a probability of 0.2 that Distribution 4 obtains. Distribution 2 has a probability of 0 because the theory is only concerned with the probability of a distribution obtaining under conditions of full compliance. Given that fully compliant people would equalize holdings (as neither party has chosen sanctionably prior to the motorcyclist's choice), it is assumed that the passerby transfers 10 units of advantage to the assisted motorcyclist, as this is what the latter is owed as a matter of justice. Thus, Distribution 2 is assigned a probability of 0, and whatever probability it would have been assigned assuming actual compliance (say, 0.1) is added to the probability of Distribution 3 obtaining assuming

actual compliance (again, 0.1) to yield the probability of Distribution 3 obtaining under conditions of *full* compliance (0.2).

The fourth row lists the probability-adjusted distributed advantage value of each distribution. The values in this row are determined by multiplying the advantage value of each distribution by its conditional probability. For example, Distribution 3 has an advantage value of 980, which is then multiplied by the conditional probability of .1 to get an expected distributed advantage value of 98. The WEDA value for the choice not to wear a helmet is then the sum of the distributed advantage values of all possible distributions, which, in this case, equals 970.

The WEDA value of not wearing a helmet can be similarly represented by the following table:

Table 6.2 *Calculating the WEDA Value of Motorcycling without Wearing a Helmet*

	Distribution 5	Distribution 6	Distribution 7	Distribution 8
	Does Not Wear a Helmet			
	Does Not Crash		Crashes	
	Transfer	Assist (No Transfer)	Assist (Transfer)	Failure to Assist
Distributed Advantage	M = 510, P = 510	M = 460, P = 500	M = 480, P = 480	M = 100, P = 500
Advantage Value	1020	960	960	600
$P(\text{Distribution} \mid \text{No Helmet})$	0.7	0	0.1	0.2
Expected Advantage	714	0	96	120
WEDA	930			

As with Table 6.1, the distributions represented in Table 6.2 are defined in terms of the sets of events that generate a particular distribution of advantage. However, Table 6.2 describes the possible distributions that might arise from the motorcyclist choosing *not* to wear a helmet. Thus, the numbers in the distributed advantage row have been adjusted to model the distributional consequences of this choice. For example, the motorcyclist is assigned more advantage in Distribution 5 than she is in Distribution 1, as it is assumed that she derives greater enjoyment from riding without a helmet. Additionally, because WEDA calculations are made under the assumption of full compliance, it is assumed that the motorcyclist transfers half of this surplus advantage

to the passerby, thereby leaving each party with 10 more units of advantage than she possesses in Distribution 1. Similarly, it is assumed that an accident without a helmet is much more severe than an accident with a helmet. Thus, the motorcyclist is assigned less advantage in Distributions 6, 7, and 8 than in counterpart Distributions 2, 3, and 4. Finally, because the fully compliant passerby makes an equalizing transfer in Distribution 7, she ends up with less advantage than she is assigned in counterpart Distribution 3.

Using the procedure described earlier, these inputs yield a WEDA value of 930 for the choice to not wear a helmet. Given that this value is 40 units less than the WEDA value of wearing a helmet, the proposed theory maintains that the motorcyclist chooses sanctionably when she declines to wear a helmet. She then forfeits a claim to a quantity of advantage that is equal to the quantity of advantage that all others would be expected to forego_{FC} absent such forfeiture. Recall that this value is calculated by dividing the total expected loss_{FC} of advantage by the total number of people and then multiplying that by the total number of people minus 1. So, in this simplified two-person world, one would divide 40 by 2 and then multiply by 1 to get 20 units of forfeited advantage. Additionally, the passerby would acquire a supplemental claim to this same quantity of advantage.

Finally, one can use these values to calculate the quantity of advantage to which each person is entitled. Specifically, one would subtract 20 units from the quantity of advantage to which the motorcyclist *would have had a claim* were her choice not sanctionable. For example, suppose that the motorcyclist crashes but is assisted by the passerby. According to Table 6.2, this pair of events leaves everyone with a total of 960 units of available advantage that can be distributed between the two parties. Were the motorcyclist's choice non-sanctionable, then a luck egalitarian principle of justice would assign her a claim to an equal quantity of advantage – that is, 480 units – as no one has made any past sanctionable choice that would justify inequality. However, because her choice *is* sanctionable, one must subtract the forfeited 20 units of advantage to arrive at a just share of 460 units. Additionally, one would add 20 units to the passerby's counterfactual share to yield a just share of 500 units. Thus, a just distribution is realized without any supplemental transfer from the passerby to the motorcyclist. The anarchist conclusion would then assign both the motorcyclist and the passerby a claim against anyone (i.e., the other party) using unowned resources in a way that would diminish – or, in the case of the motorcyclist, further diminish – her respective share. In other words, the passerby sustains all of her prior claims against the motorcyclist using various resources despite the fact that respecting those claims will now leave the motorcyclist comparatively worse off.

6.5 Anarchism without Moral Tyranny

Sections 6.3 and 6.4 have explained how the proposed theory of sanctionable choice holds people responsible for destroying_{FC} advantage such that they cannot leave others with less_{FC}. In other words, these sections have demonstrated that the theory satisfies the moral tyranny constraint in a way that standard prudential contextualism does not. However, note that this demonstration does not show that the theory *fully* satisfies the moral tyranny constraint. To fully satisfy the constraint, a theory must not only preclude sanctionable choosers like the motorcyclist from leaving others with less_{FC} but also preclude those others (e.g., the passerby) from leaving sanctionable choosers with less_{FC}. Only then will the proposed account have a theoretical advantage over contextualism.

Fortunately, the demonstration of this point is fairly straightforward. Note that contextualism's moral tyranny problem follows from (a) its claim that the sanctionability of a choice is a function of the actual advantage the chooser foregoes as a result of that choice and (b) the fact that other agents are able to determine the quantity she forgoes *after* her choice has been made. Together, these two propositions entail that other people have the power to unilaterally render a person's choice sanctionable, thereby stripping her of a claim to advantage, which, in turn, leaves her with less_{FC}. By contrast, the theory proposed here makes sanctionable choice strictly a function of the agent's choice and the evidence already available to her (namely, the evidence about the consequences of various choices under conditions of full compliance). As a result, the agent is able to avoid choosing sanctionably, meaning that others lack the ability to *unilaterally* leave her with less_{FC}. Given that such unilaterality is a necessary condition of moral tyranny, it follows that the proposed theory fully satisfies the moral tyranny constraint in a way that contextualism does not.

Suppose, for example, that, after the motorcyclist crashes without a helmet, the passerby refuses to assist her. According to the proposed theory, this choice would also be sanctionable, as the WEDA value of not assisting (600) is lower than the WEDA value of providing assistance without transfer (960).¹⁸ Thus, according to the proposed theory, the passerby would forfeit a claim to 180 units, as this value is equal to the

¹⁸ It is assumed that the probability of each distribution obtaining is 1 conditional on the choice to not assist and assist, respectively. Thus, the WEDA value of each choice is equal to the distributed advantage value of the distribution it will bring about.

total lost_{FC} advantage (360) multiplied by $\frac{2-1}{2}$. This value must then be subtracted from what would have been the passerby's just share, were her choice not sanctionable. In this case, this counterfactual just share is equal to 320. Note that there are 600 total units of advantage available to distribute, with an equal distribution assigning 300 units to each party.¹⁹ However, given the motorcyclist's past sanctionable choice to not wear a helmet, her share has to be adjusted downward to 280 units to reflect the 20 units of advantage she forfeited in virtue of that choice. These units are then reassigned to the passerby such that her counterfactual just share is 320 rather than 300 units. One then subtracts the forfeited quantity of 180 units from this amount to yield a just share of 140 units of advantage. Finally, one reassigns these 180 units to the motorcyclist, who ends up having a just share of 460 units (with the anarchist conclusion assigning her the appropriate corresponding set of distributive claims). Crucially, this is the same quantity of advantage to which she would have been entitled had the passerby assisted her; thus, assuming future full compliance, the motorcyclist will end up with just as much advantage without assistance as she would with assistance. In other words, the passerby is unable to leave the motorcyclist with less_{FC} – a result that demonstrates that the WEDA-based anarchist conclusion satisfies the moral tyranny constraint in a way that a contextualist version of the position does not.

6.6 Amending the Theory

Before completing the argument, a quick amendment must be made to the theory to avoid an objection that might otherwise undermine its plausibility. In its present form, the proposed theory makes sanctionable choice a function of the failure to maximize WEDA *assuming full compliance*. However, in many cases, people will not actually comply with others' claims. As a result, there will be cases where the choice that maximizes full-

¹⁹ When calculating the WEDA value of the motorcyclist's choice, it was assumed for simplicity that there was no opportunity for further advantage transfer conditional on the passerby declining to assist the motorcyclist. This assumption has now been relaxed so as to illustrate how advantage is to be distributed in virtue of the passerby's choice. If one were being very precise, all of the possible distributions of advantage conditional on non-assistance should have been included in the two tables and each assigned a probability. However, given that this would have added hundreds of columns to the tables without changing the result of the WEDA calculations, these distributions were excluded from the foregoing discussion.

compliance WEDA runs contrary to the demands of justice. Consider, for example, David Estlund's Slice and Patch case:

Slice and Patch Go Golfing

Suppose that unless a patient is cut and stitched he will worsen and die (though not painfully). Surgery and stitching would save his life. If there is surgery without stitching, the death will be agonizing. Ought Slice to do the surgery? This depends, of course, on whether Patch (or someone) will be stitching up the wound. Slice and Patch are each going golfing whether the other attends to the patient or not. Does anyone act wrongly? (2020, 33)

In this case, Slice choosing to operate would maximize WEDA under conditions of full compliance, as a fully compliant Patch would stitch up the wound, thereby leaving everyone with maximal equal advantage. However, given that actual Patch *will not* stitch up the wound, it seems plausible to think that justice demands that Slice refrain from operating. (For these purposes, this can simply be stipulated.) Given these premises, the posited theory entails a seemingly unacceptable result: Slice declining to operate is both a just choice and a sanctionable choice in virtue of which she forfeits a claim to advantage.

To avoid this problem, the theory can be amended as follows. Rather than define sanctionable choice strictly in terms of a failure to maximize WEDA under conditions of full compliance, a second necessary condition of sanctionable choice can be added to the theory: A person chooses sanctionably if and only if she fails to maximize WEDA under conditions of full compliance *and* fails to maximize the chances that advantage is appropriately distributed assuming actual compliance. Thus, when Slice chooses not to operate on the patient, she does not choose sanctionably, as the added necessary condition is not met.

Further, the amended theory still satisfies the moral tyranny constraint. Admittedly, when Slice declines to operate, she leaves the patient with less_{FC} advantage than if she chose to operate. However, recall from Chapter 2 that the moral tyranny constraint is only violated when a theory enables a person to *discretionarily* leave others with less_{FC}, where a discretionary action is one that is not obligatory according to the theory in question. Given that it is obligatory that Slice not operate, the fact that this choice leaves the patient with less_{FC} advantage under the proposed theory does not entail that the theory permits moral tyranny. The amended theory thereby avoids declaring just actions sanctionable while also satisfying the moral tyranny constraint.

6.7 Additional Advantages of the Theory

In addition to resolving the moral tyranny objection, the proposed theory of sanctionable choice also allows the anarchist conclusion to avoid the other problems with luck egalitarianism presented in Section 6.1. Recall, first, Arneson's objection that luck egalitarianism delivers incorrect results in cases of costly rescue, for example, by declaring Mother Teresa to have chosen sanctionably when she gives her money to the poor (Arneson 2011a, 244; 2011b, 33–4). While this objection seems like a genuine problem for prudential contextualist luck egalitarianism, there are two reasons that the proposed theory would not entail that Mother Teresa chooses sanctionably.

First, there are certain ways of filling in the details of the case such that Mother Teresa maximizes WEDA under conditions of full compliance and, thus, does not choose sanctionably. For example, if the poor are in their position due to bad luck and Mother Teresa has a comparative advantage in transferring advantage, then her actions would be expected to leave others with more_{FC} than if she focused on generating advantage. Second, even if her actions do not maximize WEDA assuming full compliance, the poor would have distributive claims that others only use natural resources in ways that increase their advantage, thereby making her transfers obligatory. Given such a duty to transfer, Mother Teresa's choice to aid the poor does not meet the second necessary condition of sanctionable choice introduced in Section 6.6. Thus, the anarchist conclusion avoids delivering an incorrect result in this case because it does not entail that she forfeits any distributive claims.

Section 6.1 also introduced Hurley's "boring problem." Recall that this problem emerges from (a) luck egalitarianism's claim that inequality is justified if and only if those with less have chosen sanctionably and (b) the standard contextualist view that sanctionable choice is a function of whether or not a person has imprudently failed to maximize her possessed advantage. When taken together, these two claims entail that a comparative relation between levels of advantage can be justified by appealing to a responsibility relation that obtains between a person and her individual holdings. However, Hurley argues that the latter relation does not appear to justify the former: The fact that a person is responsible for having a particular quantity of advantage would not seem to justify others having *more* advantage, as the person is not responsible for this *difference*. Thus, Hurley worries that luck egalitarianism is internally incoherent (2003, 160–1).

What is now apparent is that this objection is specific to contextualist versions of luck egalitarianism, as only contextualism maintains that a person's imprudent choices – that is, those that leave her with some diminished quantity of advantage – justify a comparative inequality in advantage. By contrast, the proposed theory holds a person responsible for failing to position herself and others in a way that would allow everyone to bring about justice via compliance without anyone needlessly foregoing advantage. This seems like a much more apt *justificans* for comparative inequality than a person being responsible for her own level of advantage.²⁰ Alternatively, the comparative relation might be justified by the fact that this relation obtaining is a necessary condition of avoiding moral tyranny. Regardless of the exact justification offered, the proposed theory solves Hurley's boring problem by explicating why an individual's sanctionable choice justifies inequality despite that individual not being responsible for the inequality in question. Thus, the anarchist who takes inequality to require justification (for the reasons discussed in Section 5.7) can avoid Hurley's worry that sanctionable choice is not a suitable *justificans*.

²⁰ Gerald Lang (2015; 2021) has suggested that the boring problem might be solved by modifying the *justificandum* of sanctionable choice: Rather than have such choice justify a comparative inequality between two persons – that is, the gap between their respective levels of advantage – it would, instead, justify a gap between an agent's level of advantage and some egalitarian baseline. On this proposal, each person is assigned some baseline share of advantage. If she then ends up with either more or less advantage than her assigned baseline share, that deviation would be just if and only if the difference between her share and the baseline is attributable to her choices rather than luck (2015, 706). The thought here is that, while a person is not responsible for the fact that she has less advantage *than someone else*, she *is* responsible for the fact that she ends up with less *than her baseline share*. Thus, there is no longer a justificatory gap of the kind identified by the boring problem, as the agent's sanctionable choices make her responsible for the state of affairs that those choices are supposed to justify.

However, there are three problems with this proposal. First, as Lang notes, it is unclear how to determine each person's appropriate baseline share (714). Second, the proposal seems to fail on its own terms. Lang's suggestion is that, while all *interpersonal inequalities* qualify as luck because they depend on the choices of the better-off party in addition to the choices of the worse off, deviations from the baseline depend solely on the choices of the agent and, thus, are controlled in a way that renders these deviations non-luck (and therefore justifiable). However, deviations from the baseline would equally qualify as luck (so construed), as such deviations almost always depend on the uncontrolled choices of others. For example, the person who drops below the baseline due to losing at roulette ends up in this state only because of how forcefully the casino employee spun the wheel – a fact over which she had no control. Finally, note that Lang's proposal still assumes a prudential contextualist theory of sanctionable choice where a choice justifies a deviation from the baseline if and only if it leaves the agent worse off than she might otherwise have been. Given this assumption, Lang's baseline account is still vulnerable to the moral tyranny objection, making the proposed WEDA account a superior solution to the boring problem.

6.8 The Disadvantage Creation Account

The proposed theory of sanctionable choice is similar in certain respects to a revised version of luck egalitarianism proposed by Jens Damgaard Thaysen and Andreas Albertsen (2017). Thus, it is worth explicating their theory in some detail so as to clarify the similarities and differences between the two theories. Specifically, this section will suggest that their theory is best understood as asserting that sanctionable choice is a function of how agents' choices affect the total quantity of advantage. It will then argue that the proposed WEDA-based theory of sanctionable choice has three advantages over this interpretation of Thaysen and Albertsen's account.

Like the proposed account, Thaysen and Albertsen attempt to solve the problem of costly rescues – that is, the problem illustrated by Arneson's Mother Teresa case – by revising which choices count as sanctionable. Specifically, they posit that a choice is sanctionable if and only if it *creates* disadvantage that would not have otherwise been possessed by anyone. For example, if a villain drops a brick off of a building and it strikes someone, she creates disadvantage because she leaves someone worse off while no one would have been worse off had she chosen differently (95). By contrast, the hero who pushes someone out of the way of a falling brick and gets struck herself *distributes* disadvantage, as she merely changes who possesses disadvantage that would have obtained irrespective of her choice (96). In this way, Thaysen and Albertsen's theory avoids the implication that Mother Teresa makes a sanctionable choice when she aids the poor, as she is merely distributing disadvantage to herself rather than creating disadvantage.

To fully explicate Thaysen and Albertsen's proposal, a more precise account of disadvantage creation must be provided. Their formal statement is that an "agent is responsible for creating a (dis)advantage if, and only if, she is responsible for behaving in such a way that *somebody* was (dis)advantaged" (94). If taken literally, this statement is misleading, as it suggests that an action ϕ creates disadvantage if and only if there is a person who is left worse off in the world where ϕ occurs relative to the counterfactual world where it does not.²¹ However, this is clearly not how Thaysen and Albertsen intend their analysis to be interpreted, as such an account would entail that the hero creates disadvantage due to the fact that there is a person who is left worse off by her action (namely, herself).

²¹ This is also the natural way of precisifying Thaysen and Albertsen's later restatement of duty creation wherein they assert that such creation obtains when "nobody would be worse off if not for [the agent's] exercise of responsibility" (95).

To clarify Thaysen and Andersen's analysis of disadvantage creation, it will be helpful to consider their analysis of disadvantage *distribution*, as the former is supposed to contrast with the latter. Specifically, they posit that disadvantage distribution occurs if and only if "X, rather than Y, was (dis)advantaged" by the agent's action. However, there is some potential ambiguity in this statement that calls for additional precisification. One way of interpreting this analysis is as follows: An agent distributes disadvantage by ϕ -ing if and only if X is worse off in the world where the agent ϕ -s than in the world where she does not ϕ and Y is worse off in the latter world than she is in the former. Notably, this restatement delivers the correct results in the hero case, as the hero is worse off in the world where she pushes the beneficiary than she is in the world where she does not push, while the beneficiary is worse off in the latter than she is in the former. One could then define disadvantage creation as cases where (a) the agent's action leaves someone worse off relative to inaction and (b) the action is not an instance of disadvantage distribution.

The problem with this proposal is that the restated analysis of disadvantage distribution seems to deliver incorrect results. Consider, for example, a modified case where a villain has a small quantity of fun dropping a brick on her victim. In this case, the victim is worse off in the world where the brick is dropped than she is in the world where it is not dropped; at the same time the villain is worse off in the latter world than she is in the former (because she has less fun). Thus, the restated account would entail that the villain distributes disadvantage rather than creates it. Given that this is seemingly a paradigmatic case of disadvantage creation, this result is a *reductio* of the proposed restatement.

So what is a better statement? The apparent solution is to put things in terms of the total quantity of disadvantage resulting from an action: An agent distributes disadvantage by ϕ -ing if and only if (a) some person has either more or less advantage in the world where the agent ϕ -s than the counterfactual world where the agent does not ϕ and (b) there is the same total quantity of disadvantage in the former world as there is in the latter. This account seems to deliver the correct results in the paradigmatic cases. For example, when the hero saves the beneficiary from the falling brick, the resulting world contains the same quantity of disadvantage as the world where no rescue occurred; the only difference is that the beneficiary has more advantage in the rescue world while the hero has less. Thus, the hero merely distributes disadvantage in that case. By contrast, the villain who derives enjoyment from dropping a brick off of a building does not distribute disadvantage, as there is more disadvantage in the world where

she drops the brick than there is in the world where she does not act in this way.

One final adjustment is needed. The fact that Thaysen and Albertsen use the term “(dis)advantage” when articulating their distribution/creation distinction suggests that they actually take there to be four distinct phenomena: advantage distribution, disadvantage distribution, advantage creation, and disadvantage creation. This fourfold division is incompatible with the just-posed restatement, as this restatement would only allow for disadvantage creation/distribution with there being no apparent acts of advantage creation/distribution. To fix this problem, the analysis can be amended as follows. An agent distributes *disadvantage* by ϕ -ing if and only if (a) there is the same total quantity of disadvantage in the world where she ϕ -s as there is in the counterfactual world where she does not ϕ and (b) there is less total advantage after she ϕ -s than there was just prior to her ϕ -ing. This joint counterfactual and trans-temporal comparison seems to capture the idea that people are worse off – that is, there was *disadvantage* generated – but the agent merely distributes that worsening without contributing to it. The account of *advantage* distribution would then be identical to the just-proposed analysis except that Condition (b) asserts that there is *more* total advantage after the agent ϕ -s.

The restatement of (dis)advantage creation is a bit more straightforward, as one can capture the idea of leaving everyone (worse off) better off without having to make any trans-temporal comparisons. Specifically, an agent creates disadvantage by ϕ -ing if and only if there is less total advantage in the world where she ϕ -s than there is in the counterfactual world where she does not ϕ . And she creates advantage by ϕ -ing if and only if there is more total advantage in the world where she ϕ -s than in the counterfactual world where she does not ϕ .

This explication helps to reveal the similarities and differences between Thaysen and Albertsen’s theory of sanctionable choice and the one posited by this chapter. The primary similarity is that both theories reject contextualist theories of sanctionable choice and, instead, make sanctionable choice a function – at least in part – of what effect the agent’s choice has on the total quantity of advantage (more on the qualifier later). This allows both theories to sidestep Arneson’s objection that luck egalitarianism unacceptably entails that costly rescues are sanctionable. Given that such rescues do not paradigmatically affect the total quantity of advantage, they would not count as sanctionable choices under either theory.

However, there are three important differences that give the posited WEDA-based account a theoretical advantage over Thaysen and Albertsen’s

proposal. First, their theory makes sanctionable choice strictly a function of created disadvantage rather than the *expected advantage value* of choices. As a result, their theory entails that a person chooses sanctionably when she makes a choice that maximizes expected advantage but ultimately creates disadvantage due to bad luck. For example, suppose that a person reasonably believes that there is a probability of .9 that she will create 100 units of advantage if she ϕ -s and a probability of .1 that she will create 10 units of disadvantage (i.e., -10 units of advantage). By contrast, if she does not ϕ , she will create 5 units of disadvantage with a probability of 1. Given that the expected value of ϕ -ing is 89 while the expected value of not ϕ -ing is -5, the agent chooses to ϕ ; however, she gets unlucky and generates 10 units of disadvantage. Given that there is more total advantage in the world where the agent ϕ -s than the counterfactual world where she does not, the posited restatement of Thaysen and Albertsen's proposal entails that she has created disadvantage and can thereby be held responsible – a seemingly unacceptable result.²² By contrast, the WEDA account avoids this implication by making sanctionable choice a function of expected total advantage rather than counterfactual advantage comparisons.

A second important difference is that Thaysen and Albertsen's account seemingly declares inequality to be just both when someone chooses sanctionably – that is, creates disadvantage – and *also* when someone makes a choice that creates advantage. For example, they hold that the miner who happens to strike a vein of gold that no one else would have found creates advantage and is, thus, entitled to keep some of the profits even if that results in inequality (98).²³ However, it seems inappropriate for a luck egalitarian theory to declare such a luck-based inequality just. After all, the fact that the miner was lucky enough to be uniquely positioned to extract the gold does not seem to justify her ending up better off than everyone else. By contrast, the WEDA theory incorporated into the anarchist conclusion does not posit such a category of *rewardable choice*; rather, anything short of maximizing WEDA is sanctionable while only the maximizing choice entitles a person to an equal share of advantage

²² Thaysen and Albertsen do specify that a choice is sanctionable only if it was *foreseeable* that it would create disadvantage (100). However, they do not consider cases where the disadvantage creation was foreseeable but not the reasonably expected outcome.

²³ Note the caveat that the miner is only entitled to *some* of the profits. This is because, according to Thaysen and Albertsen, she is only responsible for generating part of the created advantage, as some of that advantage is attributable merely to the resources rather than anything the miner did (98). Unfortunately, they do not provide an account explicating how one determines the portion of created advantage for which a person is responsible.

relative to the shares of other successful maximizers. Thus, the proposed theory would not license the inequality in question.

Finally, while Thaysen and Albertsen make sanctionable choice a function of counterfactual differences in total advantage, the posited account makes sanctionable choice a function of both the quantity *and* distribution of advantage in the relevant counterfactual worlds. This allows the posited WEDA account to avoid seemingly counterintuitive implications of Thaysen and Albertsen's proposal. Consider, for example, a case where *P* has a choice between realizing world *E* where she and *Q* each have 10 units of advantage or world *U* where she has 5 units of advantage and *Q* has 20 (assume, for simplicity, that the warranted probability of each outcome is 1). On Thaysen and Albertsen's account – at least, as it has been interpreted here – *P* would create disadvantage if she realizes *E*, as it has less total advantage than *U*. In other words, if *P* were to realize the egalitarian distribution, she would thereby make a sanctionable choice for which she could be held responsible. However, this result is seemingly a *reductio* of any posited theory of luck egalitarianism. By contrast, the theory proposed in this chapter would assign a higher WEDA value to *E*, thereby making *P*'s choice to realize an egalitarian distribution non-sanctionable.

In sum, Thaysen and Albertsen make the right kind of theoretical move by rejecting contextualism in favor of an account that makes sanctionable choice a function of total advantage. However, their failure to build expected value into their theory, their endorsement of rewardable choice, and their neglect of distributive considerations all compromise the extensional adequacy of their account. Thus, luck egalitarians troubled by Arneson's objection ought to adopt the proposed WEDA-based theory rather than Thaysen and Albertsen's proposal.

6.9 Conclusion

This concludes the argument for the anarchist position. It began with a fairly simple and plausible constraint on which theories of duties are acceptable. It then argued that a number of influential libertarian and egalitarian principles follow from this constraint, namely the consent theory of legitimacy, the Lockean proviso, and luck egalitarianism's incorporation of responsibility. These principles, in turn, were shown to entail other components of the anarchist position: Both the Lockean proviso and the consent theory of legitimacy independently entail the absence of external private property, while the former entails that persons can easily appropriate their bodies, thereby allowing for (near) universal

self-ownership of the kind articulated by ASO. Finally, it was argued that, in the absence of private property, both egalitarians and libertarians have reason to accept the anarchist conclusion – that is, the contention that each person has a claim against others using unowned resources in a way that (a) would leave her comparatively worse off where (b) that inequality did not correspond to any sanctionable choice on her part. This conclusion, it was argued, is both compatible with libertarian entitlement theories of justice and follows from libertarians' egalitarian approach to assigning persons rights and powers.

The purpose of this chapter was to render the anarchist conclusion fully determinate by specifying which choices qualify as sanctionable – and, more specifically, to do so in a way that brings the position into full compliance with the moral tyranny constraint. This chapter posited that a choice is sanctionable if and only if it fails to maximize warranted expected distributed advantage under conditions of full compliance (and fails to maximize the chances that advantage is appropriately distributed assuming actual compliance). Such an account ensures that no person is able to unilaterally, discretionarily, and foreseeably leave others with less_{FC}, thereby satisfying the moral tyranny constraint. Additionally, this theory of sanctionable choice allows the anarchist conclusion to avoid some of the major objections that plague standard luck egalitarian theories of distributive justice while still delivering equally egalitarian prescriptions vis-à-vis the use of natural resources.

In this way, the foregoing chapters have aimed to defend a heterodox philosophical position that synthesizes both libertarian moral principles and an egalitarian principle typically associated with the socialist left. Of course, the suggestion that these principles might be combined in this way will be intuitive to social anarchists, as their movement is composed of people who endorse (or would, upon reflection, endorse) just such a set of principles. However, they might still have been surprised to discover that this position can be largely derived from a single, simple theoretical *desideratum* and, thus, has the kind of coherence discussed in Section I.2. Similarly, libertarians may have been surprised to find that some of their core principles commit them to rejecting private property in favor of egalitarianism. And, for those who were not already sympathetic to either libertarianism or the anarchist position, the foregoing argument has hopefully demonstrated that there is at least a plausible and coherent variety of anarchism that deserves serious consideration when assessing what duties we have vis-à-vis resources and the state.